

DARKTRACE

Darktrace Offering Product Specification

020 7930 1350

Maurice Wilkes Building, Cowley Rd, Milton, Cambridge CB4 0DS

darktrace.com

Preface

This document is the Darktrace Offering Product Specification for across the Darktrace Product range and should be read in conjunction with the Darktrace Master Service Agreement and the applicable Product Order Form. In the event of a lacuna for a given Product, the relevant technical documentation as hosted in the Customer Portal applies. Service Definitions for Darktrace Additional Services may be found at <https://darktrace.com/legal/product-specifications-and-service-definitions>.

Product Specifications Table of Contents

Darktrace / NETWORK.....	3
Darktrace / OT.....	41
Darktrace / IDENTITY	51
Darktrace / EMAIL	71
Darktrace / EMAIL – DMARC.....	93
Darktrace / CLOUD.....	101
Darktrace / Forensic Acquisition & Investigation.....	132
Darktrace / Secure AI	147
Product Agnostic Details	161

Darktrace / NETWORK

1.	Introduction	3
2.	Real-Time Detection	4
2.1.	Product Features	4
2.2.	Deployment Architecture	9
2.3.	Implementation	11
2.4.	Administration	14
3.	Autonomous Response	15
3.1.	Product Features	15
3.2.	Deployment Architecture	26
3.3.	Implementation	27
3.4.	Administration	31
4.	Requirements	32
4.1.	Physical Install Requirements	32
4.2.	Darktrace Cloud Region Availability & Security	34
5.	Customer Responsibilities	36
6.	Considerations	39
7.	Roles	40

1. Introduction

Darktrace / NETWORK is a market leading cyber security solution that uses AI technology to build a dynamic understanding of customers' organizations. It models and evaluates network activity to gain visibility of threats in real-time and to respond autonomously and at machine speed.

Darktrace / NETWORK is comprised of two elements: Real-Time Detection and Autonomous Response. This Product Specification outlines the anticipated operation, architecture, implementation, and administration for each element.

This document it to be read in conjunction with the Darktrace Master Service Agreement which governs the usage of the Darktrace Product Offering. References to "Customer" throughout this document should be read to refer to the entity that is the owner of the software subscription and is ultimately responsible for its operation, whether as end-user or service provider.

Darktrace / NETWORK Subscriptions are sold according to the Usage Metrics bands set out in the relevant Product Order Form.

2. Real-Time Detection

Darktrace / NETWORK Real-Time Detection was previously named DETECT/Network and Enterprise Immune System. Darktrace / NETWORK that is implemented into cloud-hosted environments may be referred to in exceptional circumstances as Darktrace / CLOUD (Network). This specification applies for Customers who have purchased the product under any of the above naming conventions.

2.1. Product Features

2.1.1.High Level Summary

Darktrace / NETWORK Real-Time Detection evaluates network activity of devices for behavior outside of a “normal” state, created by an ongoing analysis of connections and other supplied data events across the enterprise network environment. The system profiles individual devices based upon their activity, and the activity of those it deems “peers” due to similar behavioral activity (“modelling”). Alerts are raised when specific criteria, or a minimum threshold of unusual behavior is met. The output is surfaced to operators in the Darktrace Threat Visualizer interface for investigation and resolution.

2.1.2.Data Retrieval and Ingestion

Darktrace / NETWORK Real-Time Detection retrieves and processes raw network traffic. Network traffic may be delivered directly to a physical Darktrace appliance ingestion port, to a virtualized Darktrace vSensor instance, or forwarded from Darktrace host-based osSensor agents to a Darktrace vSensor.

To ensure full visibility, all network traffic should be mirrored or ingested by Darktrace in some format. It is the responsibility of the Customer to ensure that Darktrace maintains visibility over network traffic in the event of network re-architecture outside the scope of the initial deployment design. The network traffic provided to Darktrace must be of suitable quality and not contain duplication, fragmentary data, or be incomplete in any fashion (for example, unidirectional).

Darktrace / NETWORK Real-Time Detection can also receive and parse log data in syslog from external tools (this requires additional configuration). Master, Unified View, physical Probe and vSensor instances (as defined at Section 2) support syslog ingestion.

Darktrace offers a suite of Threat Intelligence and Telemetry Integrations where data may be retrieved from other compatible security tools or OSINT sources. The retrieval method for each integration is detailed in the corresponding integration documentation. These alternative inputs (“Threat Intelligence Integrations” & “Telemetry Integrations”) are configured on the Darktrace Threat Visualizer “System Config” page.

Darktrace also provides a REST API for automation and a subset of relevant data input.

2.1.3.Darktrace Analysis

Darktrace analysis is multi-stage. The following outline highlights key elements of the process but is not intended to be an exhaustive list of all analysis performed during operation of the Darktrace platform, Threat Visualizer interface, or any underlying components.

Where source data is low quality, represents only a subset of all network activity, is delayed by external factors, or otherwise is incomplete, the quality of Darktrace analysis will be severely impacted. It is therefore imperative that data ingestion is configured comprehensively and to a high standard by the

Customer. Analysis may also be impacted by overloading, such as that observed when traffic throughput significantly increases beyond scoped levels.

2.1.3.1. Darktrace Deep Packet Inspection

Darktrace Deep Packet Inspection is performed on raw ingested network traffic to produce metadata for analysis by other platform components. The list of supported L4 - L7 protocols for inspection is regularly expanded, and analysis is continuously refined.

The output of Deep Packet Inspection is passed to the Darktrace model engine, to the core Darktrace “classifier” engine, and made available for display in the user interface and any other platform components such as the Cyber AI Analyst.

2.1.3.2. Darktrace ‘Pattern of Life’ Analysis and Classification

Network events, user activity events, connection data and any other configured inputs are subject to Darktrace ‘pattern of life’ analysis. Darktrace will create individualized behavioral profiles for the network entities it observes and surface activity which is considered inconsistent with the expected norm. This baseline is derived from - but not limited to - an analysis of the behavior of the individual entity, analysis of one or more clusters created from similarly behaving entities, and many variable factors such as time or communication protocol. The ‘pattern of life’ data is continually updated in real time, and reflects the data that has been received, with a greater weighting to more recent data.

Analysis is performed in the “classifier” stage by a multitude of classifiers. This core analysis applies many approaches including unsupervised machine learning techniques such as Bayesian meta-classification, techniques derived from graph theory and network analysis such as node/graph centrality, approaches derived from statistical analysis such as spectral clustering and many other techniques. The previous examples of utilized techniques provided are for illustrative purposes and should not be considered exhaustive.

Darktrace does not offer any capability to access underlying behavioral models or classification output. Darktrace provides the ability to utilize the output of this analysis using the *models* framework. Darktrace will create alerts to indicate anomalous activity which will be inserted into the “event log” within the Darktrace Threat Visualizer interface of the corresponding device, user, or entity. These unusual activity “notices” are also consumed by Darktrace Real-Time Detection models.

2.1.3.3. Darktrace Real-Time Detection Models Framework

The Darktrace Real-Time Detection model engine is the logical framework within which events and the output of ‘pattern of life’ analysis is evaluated. Both analysis components described above submit data to the model engine for evaluation.

A model is used to define a set of conditions which, when met, will alert the system to the occurrence of a particular event or chain of anomalous behavior. The models framework leverages both the underlying ‘pattern of life’ detection and outputs from Darktrace Deep Packet Inspection, telemetry inputs, Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust modules. Output from the complex anomaly framework is available in accessible, building block format and can be combined with simple conditions and logical expressions to create tailored activity detection.

Custom models can be defined by operators of the Darktrace Threat Visualizer to meet any criteria; these are created and maintained in the Threat Visualizer “Model Editor” interface. Darktrace bears no liability for the operation or outcomes of custom models created by Customer.

Default Darktrace models are focused on 'pattern of life' anomaly detection, potentially malicious behavior and optional compliance issues. Darktrace will periodically update these standard supplied models - customers with Call-Home or Darktrace-hosted cloud instances will receive updates automatically, clients without automatic updates will receive all applicable model updates when Darktrace Threat Visualizer software is updated. The conditions for auto-update eligibility are described in the relevant product documentation.

Models may also be used to trigger actions within the Darktrace Threat Visualizer platform; the output of the Darktrace Real-Time Detection model engine is described below.

2.1.3.4. Cyber AI Analyst

Darktrace Cyber AI Analyst performs a meta-analysis upon the previous layers of analysis described. Please refer to "Cyber AI Analyst" in the Product Agnostic section below.

2.1.3.5. Cross-Capability / Cross-Coverage Area Analysis

Where other Darktrace Real-Time Detection coverage areas are deployed, cross-platform analysis is performed by multiple components. Links are created for analysis purposes between network entities and entities modeled by other Real-Time Detection components such as, for example, user entities created by Darktrace / IDENTITY module monitoring known to be associated with a given network entity, or credential entries observed by Darktrace / NETWORK monitoring also observed as part of Darktrace / ENDPOINT coverage.

Darktrace Cyber AI Analyst will also link together entities across different coverage areas in the creation of Cyber AI Analyst Incidents and Cyber AI Analyst Incidents Events. Darktrace Cyber AI Analyst may also retrieve additional contextual data during investigation from other components if deployed, such as the retrieval of associated emails from Darktrace / EMAIL.

The output from Darktrace components may be fed to Darktrace Real-Time Detection as part of supplementary Threat Intelligence; Darktrace / Attack Surface Management can provide output of malicious asset identification into the "Watched Domains" list of Darktrace Real-Time Detection to create model alerts when a network entity accesses an Attack Surface Management-identified domain.

The above provides illustrative examples of collaborative scenarios between platform components for reference but is non-exhaustive.

2.1.3.6. Processed vs Modeled

Note, a distinction exists between activity which has been "processed" and that which has been "modeled". Data may be processed and analyzed by Darktrace Deep Packet Inspection, but not further submitted to 'pattern of life' analysis or any other analysis components. In this scenario, no "device" entity is created as a result of the event ingestion, and activity performed by the entity is not evaluated for unusual activity or by Darktrace models, except in direct relation to modeled entities.

This distinction is only relevant where specific network ranges or devices have been intentionally excluded by an operator from further analysis. The responsibility for the exclusion of these devices from Darktrace analysis, and any reduction in analysis coverage as a result, is the responsibility of the Customer as operator.

2.1.4. Output

Activity processed is surfaced in the Darktrace Threat Visualizer interface, and is optionally available through the Darktrace Mobile App, compatible output formats, or the Darktrace Threat Visualizer REST API. Darktrace / NETWORK Real-Time Detection typically outputs connection metadata from Darktrace Deep Packet Inspection, results of ongoing 'pattern of life' analysis (including notification of unusual activity), model events, Darktrace Cyber AI Analyst alerts, contextual data retrieved from configured integrations (optional), and any other log-based event data created or processed by the Darktrace system.

The platform may also output alerts about overall system health and component health.

2.1.4.1. Darktrace Deep Packet Inspection Metadata

The metadata output of Darktrace Deep Packet Inspection is displayed in the Darktrace Threat Visualizer "Advanced Search" interface. Data retention for this output is on a rolling basis and is dependent upon a number of factors such as hardware capability, traffic makeup and other operational components which contribute entries to Advanced Search. Most deployments can expect around 30 days retention. Darktrace Deep Packet Inspection metadata can be exported to external tooling by supported methods for longer retention.

2.1.4.2. Alerts

Darktrace / NETWORK Real-Time Detection produces two primary alert types: Darktrace model alerts and Darktrace Cyber AI Analyst Incidents. Darktrace model alerts and Darktrace Cyber AI Analyst Incidents can be investigated in the Threat Visualizer interface, investigated in the Darktrace Mobile App, or exported to a compatible alert consumer. Alerts are categorized for priority and filterable by multiple factors, allowing for customization of alerts displayed in Darktrace interfaces and those exported to external tools.

A secondary alert type - Darktrace System Status alerts - may also be generated where the Darktrace platform is experiencing degraded service due to health issues, invalid input, or where it is necessary to highlight changes to system administrators. Alerts include details of the originating host, the severity of the event, and relevant links to investigate or resolve the issue. Notifications are sent when a system event becomes active and can optionally be sent on resolution.

Darktrace supports alert export in both industry-standard generic forms such as Syslog or Email and custom integrations with common tools such as Splunk and ServiceNow. Alert outputs ("Workflow Integrations") are configured on the Darktrace Threat Visualizer "System Config" page. The details included in each external output may vary due to third-party restrictions on content length or supported formats.

2.1.4.2.1. Model Alerts

Darktrace model alerts are created as a result of Darktrace Real-Time Detection models; when conditions for a model are met, a model alert can be created in addition to other possible model actions. Darktrace model alerts contain details of the conditions that were satisfied in order to trigger the alert, the entity which met those conditions (for example, a device, or another model alert) and a description with recommended action points. Alerts will only be generated from ingested and inspected packets.

2.1.4.2.2. Darktrace Cyber AI Analyst Incident Alerts

Darktrace Cyber AI Analyst Incident alerts are created when Darktrace Cyber AI Analyst identifies activity considered significant enough to highlight to operators. Darktrace Cyber AI Analyst Incidents may contain a single “Cyber AI Analyst Incident Event”, or multiple linked findings aggregated together. External alerts are created when a new “Cyber AI Analyst Incident Event” is created, which may be associated with an existing “Cyber AI Analyst Incident”, or form an independent, new “Cyber AI Analyst Incident”. In the former case, mechanisms to identify the relationship between Cyber AI Analyst Incident Events are provided in the output.

In the Threat Visualizer and Darktrace Mobile App, the output of Darktrace Cyber AI Analyst analysis is aggregated on a per-Cyber AI Analyst Incident basis. New Cyber AI Analyst Incident Events added to existing Incidents will not produce independent alerts in these interfaces and are instead displayed as part of the existing Incident. Darktrace Cyber AI Analyst Incidents in these interfaces contains details of the activity identified, the Darktrace model alerts that triggered the initial Cyber AI Analyst investigation, the entities which performed the unusual activity, (for example, devices), the investigation steps Cyber AI Analyst performed, why activity was linked together (if multiple Cyber AI Analyst Incident Events), and a human-readable summary of the finding.

2.1.4.3. Reporting

Darktrace offers both manual and automatic PDF report generation, scheduled via the Darktrace Threat Visualizer “System Config” page. Details of reporting formats offered is provided on the Darktrace Customer Portal and relevant documentation.

Darktrace reserves the right to alter the content of reports offered to align with changing product and service offering.

2.1.5. Interface

The primary user interface of the Darktrace platform is the Threat Visualizer. The Threat Visualizer interface provides access to all major Darktrace user interfaces, consoles and product views - it contains both investigation interfaces and administration interfaces. A list of the user interfaces that comprise the Threat Visualizer - and detailed information on how to operate these interfaces - is provided in the relevant technical documentation.

Darktrace / NETWORK Real-Time Detection alerts are surfaced in the main Threat Visualizer dashboard which includes user customizable and filterable alert trays, a visualization of ongoing device activity over time, and access to the output of Darktrace analysis in both log and summary format. Alternative investigation views, access to review detailed metadata, report generation, and the ability to edit and create custom models (as described above) are also provided within the interface. Operators can also review and configure the devices and subnets observed by Darktrace / NETWORK Real-Time Detection, configure system settings, deploy integrations, review system health information, perform user management and other administrative tasks.

Physical Darktrace appliances and Darktrace vSensors maintain an administration console accessible only via SSH/22 (“Darktrace Console”). This console is used for advanced configuration tasks and is not available in Darktrace-hosted cloud environments.

2.1.6. Reporting

The Threat Visualizer “Audit Log” records changes made by operators such as model alert acknowledgement; the audit log can be exported via Syslog for extended retention.

The metadata output of Darktrace Deep Packet Inspection is displayed in the Darktrace Threat Visualizer “Advanced Search” interface. Combined output from this metadata, from Darktrace analysis, and from any actions performed by the platform automatically (such as tagging as a result of a model) are combined into logs which are displayed for each device. Data retention for metadata output and combined log data is on a rolling basis and is dependent upon a number of factors such as hardware capability, traffic makeup and other operational components which contribute data to the platform. Most deployments can expect around 30 days retention of Darktrace Deep Packet Inspection metadata and general device activity data. Automatic removal of older event log data does not affect the storage or training of the machine learning ‘pattern of life’ data.

Devices observed by Darktrace / NETWORK Real-Time Detection to be recently active in the network environment are detailed on the Threat Visualizer “Device Admin” page. Subnets observed by Darktrace / NETWORK Real-Time Detection in network connectivity are displayed on the Threat Visualizer “Subnet Admin” page. Metrics regarding data and event throughput are rendered on the Threat Visualizer “System Status” page.

Darktrace also provides a REST API for automated retrieval of a subset of data output.

2.2. Deployment Architecture

The appropriate combination of deployment components will vary depending on the customer network environment, for example, the balance of physical and virtualized infrastructure, or the physical location of covered networks. The appropriate deployment scenario may also change during the lifetime of the Darktrace deployment if significant changes are made to the network architecture. It is the responsibility of the client to ensure that Darktrace maintains visibility over network traffic in the event of network re-architecture outside the scope of the initial deployment design.

Each deployment requires a minimum of one “*Master*” instance to provide the capabilities outlined above. Deployments should be structured so that a single instance - Master or Unified View - sits at the top of the deployment topology to operate the Threat Visualizer interface and other relevant components.

Darktrace instances may operate in different roles within a distributed deployment; Darktrace vSensors can process raw network traffic ingested directly from virtualized networking equipment, or mirrored to it by connected osSensor agents, and forward the resultant metadata onward to a connected Darktrace instance. Darktrace instances may operate as *Probes* processing traffic and forwarding directly for analysis & storage on a *Master* or may store some processed traffic on-instance. Master instances may also directly ingest, process and analyze data.

The following outlines common Darktrace deployment architecture components. *Topology Components* are types of Darktrace hardware or software that may be deployed as part of a Darktrace deployment. *Topology Roles* are particular purposes that instances may perform within the deployment that restrict their operation. There are five roles: “Master”, “Probe”, “Unified View”, “Subordinate Master” (or “SubMaster”), and “Agent”.

Some components are limited to specific roles only; others are able to operate in multiple roles.

2.2.1. “Physical Appliance” (Topology Component)

Darktrace offers a physical Darktrace appliance for installation in a datacenter or other compatible location. Darktrace appliances are highly tuned, high-performance pieces of hardware that host the Darktrace platform. There are multiple types of Darktrace appliance, with different throughput capacities and options for data ingestion. The appropriate model will depend upon the installation location, traffic volume, traffic composition and deployment role. Darktrace appliances can only run the Darktrace platform and no other software.

Customer has responsibility for managing software updates, backups, and other system health factors for Darktrace appliances. Master

2.2.2. “Cloud Instance” (Topology Component)

Darktrace provides virtualized Darktrace deployments by hosting a cloud-based instance within Darktrace cloud environments, hosted by AWS and Microsoft Azure. Virtualized deployments receive data from local Probes in the customer network. Cloud instances are managed and maintained by Darktrace; software updates, backups and instance scaling are managed by Darktrace operations.

Cloud instances can operate as “Master” and “Unified View” roles but cannot operate as Probes.

2.2.3. Master (Topology Role)

A Darktrace Master instance operates analysis, modeling and alerting capabilities and serves the Darktrace Threat Visualizer user interface. Darktrace Master instances can ingest and process network traffic directly, but this is not a necessity if this action performed by other instances in the topology. A Master instance may be a physical Darktrace appliance, or Darktrace Cloud Master, hosted by Darktrace in one of AWS or Azure in an appropriate availability zone. In the latter case (Cloud-hosted Master), Darktrace does not support direct traffic ingestion, and one or more Probe instance(s) will be necessary for network traffic retrieval.

2.2.4. Probe (Topology Role)

A Darktrace Probe receives data input such as network traffic (typically from SPAN, traffic mirroring or network TAPs) or syslog. Probes do not operate a user interface, analyze ‘pattern of life’ or have a model engine. A Probe may be a physical Darktrace appliance configured to operate in a Probe role, or a virtualized Darktrace vSensor Probe hosted by the client. Probes require a connected Master and will communicate in a method defined by the client during configuration; selection of an appropriate communication mode for the deployment scenario is the responsibility of the client.

2.2.4.1. Physical Probe (Topology Component Operating in Probe Role)

A physical Probe is a hardware appliance as described above, configured to operate as a Probe.

2.2.4.2. Virtual Probe (Topology Component with Exclusive Probe Role)

The Darktrace vSensor is a virtual Probe intended for deployment in cloud-based networks or environments where it is not feasible (such as virtualized networks), or not desired by the customer, to deploy a physical Probe. vSensor software is provided by Darktrace in various pre-packaged virtual machine formats or can be installed on virtual machine instances running a compatible Linux-based operating system. The vSensor supports autoscaling in compatible environments.

The vSensor receives data input such as network traffic or syslog but does not operate a model engine or serve a user interface other than a simple configuration console. vSensors require a connected Master to which they forward processed metadata and will communicate in a method defined by the client during configuration; selection of an appropriate communication mode for the deployment scenario, and definition of a secure access password (where relevant), is the responsibility of the client.

For usage in cloud or other environments where it is not possible to span a virtual switch, the vSensor also supports ingestion of traffic from multiple osSensors.

2.2.5. Darktrace Agents

Darktrace Agents are software components deployed on customer servers to collect traffic and enrichment data, which is forwarded to associated Probes or Masters; they do not perform any data processing themselves.

2.2.5.1. osSensor Agent (Topology Component with Exclusive Agent Role)

Darktrace osSensors are lightweight, host-based server agents. osSensors are provided for a range of compatible Windows and Linux-based operating systems. osSensors do not perform analysis or server a user interface. The osSensor forwards an unprocessed copy of network traffic to a connected vSensor for processing. Connection between the vSensor and osSensor is secured using a shared HMAC token. osSensors are not suitable for connectivity over untrusted networks.

The osSensor is also offered in a containerized format via Docker Hub for deployment in environments such as Kubernetes and AWS Fargate.

2.2.5.2. containerSensor Agent (Topology Component with Exclusive Agent Role)

The Darktrace containerSensor is a tracking sensor which enhances 'pattern of life' analysis in containerized environments. The containerSensor relays metadata about workloads to the connected Master, allowing traffic observed within the cluster to be assigned to consistent, recognizable entities.

2.2.5.3. Server Agents (Topology Component with Exclusive Agent Role)

Darktrace Server Agents collect enrichment data from customer servers which is not accessible from network traffic (typically remote access solutions), then relay this to the connected Master to improve user & device tracking. Further information can be found on the Customer Portal

2.2.6. “Unified View” (Topology Role) and “Subordinate Master” (Topology Role)

A Darktrace “Unified View” (UV) is a component of a distributed Darktrace deployment, where a single Master instance with or without Probes is not sufficient, or not suitable, for the deployment scenario. In this mode, an additional instance is deployed above, and receives data from, all Masters for display in a single user interface. Darktrace Master instances under a UV operate as subordinate Masters (“subMasters”), where a subset of capabilities is removed and controlled/operated centrally at the Unified View level. In this mode, each Master continues to operate a model engine and perform ‘pattern of life analysis’, but select components are synchronized from the Unified View (such as Darktrace models and system configuration).

2.3. Implementation

Darktrace / NETWORK Real-Time Detection is implemented through deployment of one or more physical or virtual components, including at least one Master instance. Full appliance specifications, implementation guides and administration guides are all found on the Customer Portal.

2.3.1. Physical Instances

Darktrace appliances are installed within the customer network to ingest network data into one or more packet capture ports. This information is processed by the platform and the output is displayed in the Darktrace Threat Visualizer of the associated Master.

2.3.1.1. Admin Interfaces

All Darktrace instances have a network port for the admin interface, which must be connected for: access to the User interface; connection between a Probe and its Master (dependent on the configured *role*); the Call-Home function; and additional network services such as third-party integrations, monitoring, inbound enrichment data, or outbound alerting.

Darktrace physical appliances also have an out-of-band (OOB) management interface. Darktrace strongly recommends connecting the OOB to the customer network for additional hardware-layer management, troubleshooting and maintenance capabilities, particularly when the admin interface is unavailable.

2.3.1.2. Analysis Interfaces

Darktrace instances feature different analysis interfaces for ingestion of network traffic into the Deep Packet Inspection engine. The available interfaces on each type of physical or virtual instance are detailed in appliance specifications available on the Customer Portal.

During implementation, customer must ensure a port mirroring, SPAN, TAP or other data forwarding method is set up through one or more interfaces on any Darktrace instance with the Probe role, or any Master also intended to ingest network traffic.

2.3.1.3. Data checks for new instances

To confirm instances are receiving traffic and the user interface is accessible, post installation, data checks should be carried out. Details of the necessary checks for each instance are provided in the Customer Portal documentation.

2.3.2. Darktrace Virtualized Instances

Darktrace provides virtualized Darktrace deployments by hosting a cloud-based Master instance within Darktrace cloud environments (hosted on AWS and Microsoft Azure). Within the scope of Darktrace / NETWORK, Cloud Masters receive data from local Probes in the customer network (physical or virtualized). Cloud Masters may also receive traffic as part of other Darktrace products (for example, client sensors as part of Darktrace / ENDPOINT)

Cloud Master instances can analyze and ingest the same data as physical Master instances.

Organizational virtual network traffic can be sent via a Darktrace vSensor in one of two encrypted communication modes. vSensors can be deployed as a standalone virtual machine in a traffic-mirroring scenario, or with up to 255 osSensor agents (per vSensor). Darktrace osSensors can be installed on devices running Windows, supported Linux distributions and any Linux environment running the Docker engine.

In addition to processing and transmitting network traffic, vSensors can ingest and forward syslog-format logs to the Darktrace Cloud Master.

Individual, separate instances are provisioned for each customer within the relevant cloud-provider region. Data ingested from Probes in the customer network is encrypted in transit and will not leave the region. In addition, two-factor authentication is enforced on all user accounts. Darktrace is ISO27001 certified, ensuring we maintain a high standard of information security.

2.3.2.1. Configuration and Management of Cloud Masters

Cloud Masters are managed and maintained by Darktrace operations; the management and system administration console is not available. Darktrace manages instance scaling in line with traffic-load. If there is a hybrid Unified View deployment containing both physical and cloud submasters, Darktrace operations manages the cloud Master, whereas the physical Masters are managed by the customer. If you wish to modify a setting on a cloud master or submaster or configure a process that would normally require console access, please contact Darktrace support who will assist you with this process.

Threat Visualizer software on a Cloud Master is automatically updated when a new version becomes available; where possible, updates will be applied outside standard business hours. If this is not possible, the update process will cause minimal disruption for Threat Visualizer users. If automatic updates are turned off by customer request, it is customer responsibility to notify a Darktrace via the Customer Portal when they are ready to upgrade the deployment and the Darktrace team will update the deployment during the designated window provided by customer.

Multiple short-term snapshot backups of Cloud Masters are taken on a rolling basis to ensure continuity in a disaster recovery scenario.

2.3.2.2. Deployment process for Darktrace vSensors

In the same way as for physical instances, described in sections 3.1.1 and 3.1.2, vSensors need to be configured by the customer with both an admin interface and an analysis interface.

vSensors can ingest and process traffic from physical networks, in addition to virtualized traffic, with additional configuration. The vSensor supports VXLAN and ERSPAN traffics type I and type II, as well as GRE with transparent ethernet bridging.

For usage in cloud or other environments where it is not possible to capture data from a virtual switch, the vSensor also supports ingestion of traffic from native mirroring technologies, such as AWS VPC Traffic Mirroring and GCP Packet Mirroring, or multiple connected osSensors. The osSensor agent is installed on each customer device where visibility is desired and monitors all of the network traffic to/from configured interfaces of that device; the monitored traffic is then sent to the vSensor for analysis. osSensors utilize host resources to forward traffic, so should only be installed where it is not possible to retrieve traffic through other means.

2.3.2.2.1. vSensor Deployment Methods

Darktrace vSensors may be deployed via Standalone Image, Cloud QuickStart or Cloud CLI. recommended by Darktrace. The deployment guides for each method are available in the Customer Portal. It is the customer's responsibility to review the guides and decide which option works best for their organization.

2.3.2.2.2. vSensor health checks

Health checks must be performed by the customer to confirm that the vSensor is running and has connectivity with the associated master and osSensors (if applicable). Key areas to test to verify health

include: (i) that the vSensor virtual machine is running; (ii) verification of incoming packets; and checking for vSensor overload.

2.3.2.3. Deployment process for Darktrace osSensors

The Darktrace vSensor coordinates with the osSensors associated with it, ensuring traffic is captured only once when osSensor devices communicate to each other. Each osSensor registers with a vSensor using a shared HMAC token which should be supplied to both ends. It is recommended that the associated vSensor is configured in advance of osSensor setup, in order to ensure the necessary HMAC token and IP Address for the vSensor have been collected.

The osSensor is deployed as a package installed on the customer server. Each osSensor registers with a vSensor using a shared HMAC token which should be supplied to both ends. The associated vSensor should be configured in advance of osSensor setup, in order to ensure the necessary HMAC token and IP Address for the vSensor have been assigned.

The Darktrace vSensor coordinates with the osSensors associated with it, ensuring traffic is captured only once when osSensor devices communicate to each other. The installation guides for each available formats of osSensor can be found on the Customer Portal

2.3.2.3.1. osSensor Health Check

It is recommended that the Customer checks: (i) that the osSensor service is running and that there is connectivity with the associated vSensor; (ii) how many osSensors are running; and (iii) for vSensor overloading.

Further information on the testing commands can be found in Customer Portal. It is a Customer responsibility to run periodical health checks to ensure the osSensor service is still running as optimally as initially setup.

2.3.3. Architecture Configuration

After each component supporting Darktrace / NETWORK is deployed, each needs to be configured into the wider Darktrace architecture. Darktrace can provide support and advise, but configuration is a Customer responsibility to implement. For more detailed information and examples of Darktrace architectures, see the Customer Portal documentation.

2.4. Administration

General administration of the Darktrace deployment is a Customer responsibility and is performed in the Threat Visualizer interface.

The majority of configuration, including administration of system settings such as proxies, authentication configuration such as LDAP and SAML SSO, deployment of alert and threat intelligence integrations, and other administrative tasks are performed in the “System Config” page. User and group management, including the assignment of data visibility and permissions, is performed on the “Permissions Admin” page. Device administration such as the setting of labels, application of tags, and altering of priority is performed on the “Device Admin” page. Management of subnets observed by the Darktrace / NETWORK Real-Time Detection instance are administered on the “Subnet Admin” page. System health information and system alert resolution is performed on the “System Status” page.

For physical Darktrace appliances of any role, additional administrative tasks such as interface configuration, setting of host variables, and Call-Home configuration may be performed in the appliance

console, accessible over SSH. For cloud hosted instances, these administrative actions are managed by Darktrace operations directly and no access to the console is available.

Darktrace physical and cloud instances are each seeded with random passwords and two-factor authentication secrets at build time. These initial secrets are stored by Darktrace. Credentials granting access to the Darktrace Threat Visualizer interface and, for physical instances, the SSH administration console ("Darktrace Console") are provided via the Customer Portal; these passwords can be optionally changed by the client at any time. Further access administration is a Customer responsibility.

Darktrace vSensors only provide access for administration tasks via the management console. osSensor, containerSensor and other Server Agents do not provide administration interfaces and are instead managed by the customer with configuration files and/or dialogues during installation, or through the customer server they are hosted on.

Detailed information about individual administrative tasks is outlined in the relevant Customer Portal documentation for the deployment of the specified component or administrative action intended to be performed.

3. Autonomous Response

Darktrace / NETWORK Autonomous Response was previously named RESPOND/Network and Antigena. Darktrace / NETWORK that is implemented into cloud-hosted environments may be referred to in exceptional circumstances as Darktrace / CLOUD (Network). This specification applies for Customers who have purchased the product under any of the above naming conventions.

3.1. Product Features

3.1.1.High Level Summary

Darktrace / NETWORK provides Autonomous Response capabilities against network entities. Actions are triggered by specific criteria in Darktrace / NETWORK real-time detection or can be taken manually by an operator.

The Darktrace / NETWORK component applies Autonomous Response capabilities to physical and cloud network devices by controlling connectivity. This control is enacted using spoofed TCP RST packets issued by a Darktrace master, a Darktrace probe, or a virtual sensor such as a vSensor or osSensor. The actions available to Darktrace / NETWORK Autonomous Response can also be extended through integrations with endpoint security providers and with a number of popular firewalls, allowing Darktrace to enact connectivity controls using the rule or policy capabilities available in the external platform.

Please note, Darktrace / NETWORK Autonomous Response was previously known as both "Antigena" and "Darktrace RESPOND NETWORK". This terminology is preserved in many places for the purposes of ensuring continuity.

Darktrace / NETWORK Autonomous Response expands Cyber AI response to devices by severing network connections, restricting access and quarantining devices by limiting connectivity. Actions can be taken when a device exhibits significantly anomalous behavior, when it contravenes a compliance policy, when a device attempts to access a specific watched endpoint, manually taken by an operator, or as a result of any other custom criteria defined in a Darktrace / NETWORK real-time detection model.

3.1.2.Data Retrieval and Ingestion

Darktrace / NETWORK real-time detection retrieves and processes raw network traffic. Network traffic may be delivered directly to a physical Darktrace appliance ingestion port, to a virtualized Darktrace vSensor instance, or forwarded from Darktrace host-based osSensor agents to a Darktrace vSensor. Darktrace / NETWORK Autonomous Response may also process network traffic delivered to Darktrace / NETWORK real-time detection when necessary to create network-level actions.

Darktrace / NETWORK Autonomous Response uses metadata from individual network packets to create “spoofed” response packets which are then sent to participating devices. This metadata must be retrieved directly from network traffic in order to create accurate response. Darktrace / NETWORK Autonomous Response does not process all network traffic - only relevant traffic necessary for action creation. This processing and analysis are separate to the Darktrace Deep Packet Inspection performed by Darktrace / NETWORK real-time detection on raw ingested traffic. Processing is only performed for the purposes of action creation; Darktrace / NETWORK Autonomous Response does not create any metrics or surface any metadata from this analysis.

To ensure Darktrace can create action-level blocks, all network traffic should be mirrored or ingested by Darktrace in some format. It is the responsibility of the Customer to ensure that Darktrace maintains visibility over network traffic in the event of network re-architecture outside the scope of the initial deployment design. The network traffic provided to Darktrace must be of suitable quality and not contain duplication, fragmentary data, or be incomplete in any fashion (for example, unidirectional).

Darktrace / NETWORK Autonomous Response may also retrieve configuration data from third-party firewall components where required for the creation of corresponding firewall policies or rules, or to confirm that configuration was performed successfully. Darktrace is not liable for interconnectivity issues resulting from software changes or updates made to third-party components.

3.1.3.Darktrace Analysis

Darktrace / NETWORK Autonomous Response does not perform direct analysis on ingested or modeled data; actions are taken as a result of data processed and analyzed by the equivalent Darktrace / NETWORK real-time detection capability.

3.1.3.1. Darktrace Real-Time Detection Models Framework

The Darktrace Real-Time Detection model engine is the logical framework within which events and the output of Darktrace ‘pattern of life’ analysis are evaluated.

A model is used to define a set of conditions which, when met, will alert the system to the occurrence of a particular event or chain of anomalous behavior.

The models framework leverages both the underlying ‘pattern of life’ detection and outputs from Darktrace Deep Packet Inspection, telemetry inputs, Darktrace / IDENTITY, and Darktrace / CLOUD modules. Output from the complex anomaly framework is available in accessible, building block format and can be combined with simple conditions and logical expressions to create tailored activity detection.

Custom models can be defined by operators of the Darktrace Threat Visualizer to meet different criteria; these are created and maintained in the Darktrace Threat Visualizer “Model Editor” interface. Care should be taken when defining custom models to ensure that the number of alerts produced is not excessive and does not impact the system’s ability to function, or otherwise degrade service. Similarly, modification of existing default models should not result in excessive activity or alter the logic to the extent that alerts criteria can no longer be met. The responsibility to ensure that models created or edited remain within these reasonable expected boundaries lies with the Customer.

Default Darktrace models are focused on 'pattern of life' anomaly detection, potentially malicious behavior and optional compliance issues. Darktrace will periodically update these standard supplied models - customers with Call-Home or Darktrace-hosted cloud instances will receive updates automatically, Customers without automatic updates will receive all applicable model updates when Darktrace ActiveAI Security Platform software is updated. The conditions for auto-update eligibility are described in the relevant product documentation.

3.1.3.2. Autonomous Response Models

In default operation, Darktrace / NETWORK Autonomous Response responses are triggered by model alerts from a specific subset of Darktrace models, categorized as Autonomous Response ("Antigena") models.

Darktrace Autonomous Response models may directly look for specific behavior, or for indicators identified by other models operating within the real-time Detection environment. Darktrace / NETWORK models fall typically into the second category - "meta-models" - which are triggered by an alert of another Darktrace Real-Time Detection model.

Any modification to an underlying real-time detection model which results in increased or reduced model alerts will subsequently impact Darktrace / NETWORK Autonomous Response models and actions. Any under-activity or over-activity of the corresponding Darktrace / NETWORK Autonomous Response model as a result of this type of modification is the responsibility of the Customer.

Darktrace provides the ability to limit actions created by these models to a minimum score threshold.

3.1.3.2.1. Darktrace / NETWORK Autonomous Response Models

When Darktrace / NETWORK Autonomous Response is enabled within a Darktrace environment, a subset of additional models become available. These models are categorized into four high-level categories, where each category contains Darktrace / NETWORK Autonomous Response models which target specific activity within that category: "Compliance", "External Threat", "Insider Threat" and "Significant Anomaly".

Each model is intended to trigger on specific types of connection or activity and will perform different Autonomous Response actions depending on the activity identified.

Darktrace / NETWORK Autonomous Response models within these four pre-defined categories contain logical criteria that will only evaluate successfully if the device entity that has met the criteria is also tagged with the corresponding Autonomous Response tag for the given model category. This concept is referred to as "Tag-Based Eligibility" and is outlined below.

The contents of each model folder are subject to change as Darktrace personnel regularly revise and update categorized models. Darktrace may also provide models within these categories which are disabled by default to allow operators to perform customization and tuning, or where detection may only be relevant to specific compliance scenarios (optional). Models of this type can be enabled at any time by the Customer.

3.1.3.3. Darktrace / NETWORK Autonomous Response Success

Darktrace / NETWORK Autonomous Response may analyze network traffic ingested by Darktrace / NETWORK real-time detection to establish whether actions were sufficient to end the targeted connectivity.

This analysis identifies connections which have continued to transmit data after an RST packet was issued by one or more components of the Darktrace ActiveAI Security platform. The exact calculations and thresholds used to determine this success state remain subject to change.

3.1.4. Output

Darktrace / NETWORK Autonomous Response creates automatic, network-level responses to anomalous behavior detected by Darktrace / NETWORK real-time detection or when triggered manually by a Darktrace operator.

3.1.4.1.1. *Action Methodology*

There are two methods that Darktrace / NETWORK Autonomous Response may invoke to control connectivity when an action is created - spoofed TCP RST packets sent directly from a Darktrace instance, or policy/rule-based actions through third-party network components. The action methodology is selected at the time of action creation. Methods are not mutually exclusive; multiple methods may be invoked as the result of a single action.

3.1.4.1.1.1. *TCP RST Actions*

Standard Darktrace / NETWORK Autonomous Response actions utilize “spoofed” TCP RST packets to trigger the source and/or destination device to end the connection. These TCP RST packets may be issued by any compatible Darktrace instance including physical Darktrace appliances, virtualized probes such as vSensors, or on-host osSensors.

Darktrace / NETWORK Autonomous Response takes information from ingested network traffic to create packets imitating the ongoing connection, which are then sent to each end of the target connectivity imitating the other. Darktrace / NETWORK Autonomous Response will send a number of packets to each device with varying TCP and IP header properties to increase likelihood of successful reception by the endpoint (and therefore successful interruption of the TCP connection).

In internal-to-internal connection scenarios, spoofed TCP RST packets will be directed at both source and destination device - only one of the connected devices needs to receive an RST packet for the connection to be terminated. Where one device is external (internal-to-external, external-to-internal), packets will be targeted at the internal connection device.

In specific scenarios where the source or destination are unreachable, it may be necessary to target an intermediary device such as a proxy server.

The ability for Darktrace / NETWORK Autonomous Response RST packets to reach the target connection is referred to as “reachability”. Darktrace offers multiple implementation modes to ensure reachability across the network environment including dedicated “firing” interfaces connected to different areas of the network, configuration for custom packet traversal routes, and VLAN tag support. This concept is outlined in greater detail below under the administration section.

3.1.4.1.1.2. *Integration Actions*

Darktrace / NETWORK Autonomous Response extends control to non-TCP network traffic by integrating with a selection of compatible firewalls and networking components. These integrations can be configured and used independently of - although typically in conjunction with - Darktrace / NETWORK Autonomous Response's TCP RST capabilities.

Utilizing a Darktrace / NETWORK Autonomous Response firewall integration offers an alternative route to neutralize connections crossing the network boundary, where UDP traffic is frequently seen, or in scenarios where Darktrace / NETWORK Autonomous Response RST packets may not be able to traverse the network from the Darktrace instance to their intended destination.

Specific implementation varies between third-party product model and vendor - for example. Darktrace may programmatically access the networking component to create rules/policies corresponding to the desired connection-level blocks, or the network component may actively poll Darktrace for endpoints intended for blocking.

Actions created via integrations are mapped as closely as possible to the desired inhibitor (see below). If a firewall does not offer the highly specific, targeted action desired by Darktrace / NETWORK Autonomous Response, an alternative, less specific action may be taken.

3.1.4.2. Darktrace / NETWORK Autonomous Response Action States

Darktrace / NETWORK Autonomous Response actions can be taken when: Darktrace identifies significantly anomalous device behavior; Darktrace identifies activity which contravenes a compliance policy; a device attempts to access a specific watched endpoint; manually triggered by an operator; or as a result of any other custom criteria defined in a Darktrace model.

If the current system configuration permits autonomous actions - or the action was triggered manually - the action will be created in *active* state. This created action will appear on the Response Actions Page “Network” tab under “Active” actions, in the Darktrace Mobile App, and be accessible via the Darktrace API. The device will be visually distinguished (at present, with a green highlight) across the Darktrace platform for the duration of the action to indicate it is currently undergoing a Darktrace / NETWORK Autonomous Response action. If alerts for Autonomous Response actions (“*created*” state) are configured, an alert will also be sent to configured outputs.

An *active* Darktrace / NETWORK Autonomous Response action permits Darktrace / NETWORK to take the responses defined by the action configuration. For example, to send corresponding TCP RST packets automatically when matching activity is seen or make any necessary configuration changes in third-party environments such as firewalls to control connectivity. As manual Darktrace / NETWORK response actions are created directly by an operator, actions are always created in the *active* state.

If Darktrace / NETWORK Autonomous Response is compelled to request human confirmation due to the current system configuration, the action will be created in *pending* mode. This created action will instead appear on the Response Actions Page “Network” tab under “Pending” actions until a user confirms that Darktrace / NETWORK Autonomous Response can take action. A notification requesting that a human user approve the action is sent to the Darktrace Mobile App and any configured Darktrace Autonomous Response alert outputs (“*pending*” state). The device will also be visually distinguished (at present, with an orange highlight) to indicate human approval is pending for a Darktrace / NETWORK Autonomous Response action. Once confirmed by an operator, the action will move to “Active” as before.

Pending actions are those which Darktrace / NETWORK Autonomous Response wishes to take but must require approval before it can act automatically; pending actions may therefore be referred to as “human confirmation mode” actions. Pending actions do not create firewall rules (or equivalent) or cause the sending of TCP RST packets until approval is granted in some format, at which the point the action becomes “active”.

In the event that an Autonomous Response action is delayed due to the system configuration compelling the action to be created in pending mode, Darktrace bears no liability for any resulting connections, actions or damage that occurs as a result of the delay.

Darktrace provides mechanisms to alert users on the creation of pending actions requiring human approval to proceed. These notifications are surfaced in the Threat Visualizer interface, in the Darktrace Mobile App, and in compatible alert outputs.

3.1.4.2.1. *Additional Action States*

Darktrace Autonomous Response actions proceed through several states during their lifecycle. For example, an action is automatically created that *requires confirmation* (“pending”), it is then *confirmed* by a user. Another user *extended* the action, then *cleared* it. Finally, it was *reactivated* by another user, before *expiring*.

There are four key states (“pending”, “active”, “cleared”, “expired”) and two additional states (“extended”, “reactivated”):

- The first two states - pending and active are outlined above.
- Cleared actions are those which have been manually ended by an operator. Clearing informs Darktrace to cease the action, and to suppress the combination of Darktrace / NETWORK Autonomous Response action and model alert conditions for the remainder of the action’s set duration.
- Expired actions are those which are historic, regardless of their state (pending, active) before the action period passed.
- Reactivated actions are those which were made active again by an operator.
- Extended actions are those which are currently active and have had their duration manually extended by an operator.

Actions can be extended, reactivated or cleared from the Response Actions page, the Darktrace Mobile App, and the Darktrace Threat Visualizer API.

3.1.4.3. *Darktrace / NETWORK Autonomous Response Actions in Default Operating Mode*

Default operating mode (*default operation*) is here defined as Darktrace / NETWORK Autonomous Response operating in a state where autonomous (*active*) actions can be taken automatically against eligible devices when triggered by a default Darktrace / NETWORK Autonomous Response model.

Default operation presumes that Darktrace / NETWORK Autonomous Response is licensed and enabled globally on the Darktrace ActiveAI Security Platform instance, and Darktrace / NETWORK Autonomous Response has been configured according to the requirements of the network architecture and tested to ensure actions can be routed to the desired destination.

The conditions for the creation of action in any state in default operation are:

1. A Darktrace / NETWORK Autonomous Response model which has its logical criteria sufficiently met.
2. The prior criteria are met by a device or entity with the appropriate eligibility for the given model category.
3. A Darktrace / NETWORK Autonomous Response action (inhibitor) configured within the triggered model to define the action taken.
4. A system operating mode (schedule) which does not prevent the creation of a Darktrace Autonomous Response action at the given time.

The conditions for an *active* action in default operation are:

1. A Darktrace / NETWORK Autonomous Response model which has its logical criteria sufficiently met.
2. The prior criteria being met by a device or entity with the appropriate eligibility for the given model category.
3. An appropriate action (inhibitor) applied to the triggered model to define the action taken.
4. The model in question is configured to *permit* or *force* autonomous actions.
5. The system operating mode permits the creation of an active Autonomous Response action at the given time for the given device.
6. The model criteria are met by a device which does not possess configuration preventing autonomous actions.
7. The system operating mode permits the creation of an active Autonomous Response action at the given time for the given device.

Manual actions triggered by an operator are not subject to the conditions outlined above; these actions are always active on creation.

3.1.4.4. Manual Darktrace / NETWORK Autonomous Response Actions

Darktrace provides mechanisms to create Darktrace / NETWORK response actions manually. These actions can be created by a user with appropriate permissions from the Darktrace ActiveAI Security Platform Threat Visualizer interface, from the Darktrace Mobile App, or via the Threat Visualizer API.

Darktrace provides the option to enforce a free-text justification which must be submitted by the end-user when the action is created.

Manual response actions require Darktrace / NETWORK Autonomous Response to be enabled globally. Actions are automatically created in an “active” state and do not require human activation. All devices are considered eligible for manual Darktrace response actions - the responsibility to ensure that the targeted device is appropriate lies with the Customer.

The type of action taken (*inhibitor*) is selected during creation. The responsibility to select an appropriate inhibitor is also the responsibility of the manual action creator.

3.1.4.5. Model-Triggered Darktrace / NETWORK Autonomous Response Actions

In default operating mode, Darktrace / NETWORK Autonomous Response actions are triggered by models from the four categories outlined above.

This triggered action is referred to as a “model action” - a system action taken in response to a specific model criterion being met. The creation of a Darktrace Autonomous Response action is a model action which exists by default on applicable Darktrace-created models contained within the “**Antigena**” category. When the criteria for these models are met, Darktrace will invoke all model actions applied to the model, triggering a Darktrace / NETWORK Autonomous Response action as a result.

Operators may also add the Darktrace Autonomous Response model action to modified or newly created custom models. The responsibility for selecting an appropriate action inhibitor (see below), and for any actions created as a result of a custom model configuration, lies with the Customer.

3.1.4.6. Darktrace / NETWORK Autonomous Response Inhibitors

The actions Darktrace / NETWORK takes in response to a trigger will vary based upon the action configuration and the activity detected. The types of response actions taken are henceforth referred to as “inhibitors”, as they *inhibit* a specific form of behavior or connectivity.

The type of action Darktrace / NETWORK Autonomous Response takes is referred to as an inhibitor. Inhibitors can be “automatic”, where Darktrace / NETWORK selects the most appropriate action at the time, or pre-defined from a list of options.

In default operation, when the criteria for a Darktrace / NETWORK Autonomous Response model are met, Darktrace / NETWORK will attempt to create an action corresponding to the selected inhibitor for the trigger device or entity. The inhibitor is defined in the model directly.

Manual actions require the inhibitor to be selected during creation. This is applicable to actions created through the Darktrace Threat Visualizer, the Darktrace Mobile App, or the Darktrace Threat Visualizer API.

For actions created as a result of a user-created or user-modified model, Darktrace / NETWORK Autonomous Response will attempt to create an action corresponding to the inhibitor as defined in the model for the trigger device or entity.

At present, there are six pre-defined Darktrace / NETWORK inhibitors of increasing severity which may be invoked through Darktrace / NETWORK Autonomous Response models or manual response actions: “Automatic”, “Block Matching connections”, “Enforce pattern of life”, “Block all outgoing traffic”, “Block all incoming traffic” and “Quarantine device”.

Typically, Darktrace / NETWORK Autonomous Response will take the inhibitor selected in the model action in response to the Darktrace / NETWORK Autonomous Response model criteria being met. However, there are necessary exceptions to this behavior:

- Actions are proportional to threat and may be escalated if granular blocks are not sufficient.

For example, an action that targets *matching connections* is no longer sufficient where the number of connections deemed to be anomalous increases rapidly - Darktrace / NETWORK Autonomous Response will therefore escalate the severity of the action taken.
- Where it is necessary to invoke integrations with firewall vendors, the actions possible in the vendor environment may not be sufficiently granular to achieve the targeted block initially selected.

In this case, Darktrace / NETWORK Autonomous Response will modify the action to match the capabilities of the integration. The limitations of specific integrations are outlined in the appropriate integration documentation.
- Where the device or entity has a protected characteristic which excludes it from the selected inhibitor.

3.1.4.7. “Flag for Antigena” / iagn

Darktrace / NETWORK includes a “Watched Domains” component - a list of hostnames, IPs and endpoints which, if observed in network traffic or operation, will trigger an alert of some form. Entries can be added manually, through the API (*/intelfeed*), through the Darktrace / Attack Surface Management integration, or via configured STIX/TAXII intelligence feeds. No default list exists for watched domains.

When endpoints are populated by operators - programmatically or through manual entry - the option is offered to take automatic Darktrace / NETWORK Autonomous Response actions against network

connectivity to/from these endpoints. If this option is selected by an operator and an endpoint is subsequently seen in network connectivity, any Autonomous Response actions triggered are not surfaced in the user interface and are taken regardless of eligibility, device type, or how autonomously Darktrace / NETWORK Autonomous Response is permitted to act at the given time.

3.1.4.8. Autonomous Response Actions Initiated by Cyber AI Analyst

Cyber AI Analyst may initiate action creation - this functionality is optional and disabled by default. Darktrace Cyber AI Analyst can create Autonomous Response actions for detected activity which does not have corresponding active Darktrace Autonomous Response actions at the time of creation. This capability creates actions for “edges” between Darktrace Cyber AI Analyst incident events.

Actions can be created as a result of Cyber AI Analyst Incidents. Cyber AI Analyst links related devices and unusual activity together to create an overall incident structure. If Cyber AI Analyst detects that a link between two entities does not have a corresponding Autonomous Response action - and an inhibitor is available which would suitably target the activity exists - it may create an action to target this behavior.

3.1.4.9. Tag-Based Eligibility for Darktrace / NETWORK Autonomous Response

Device and credential entities can be “tagged” - marked with a specific, predefined identifier - within the Darktrace system. This may be performed manually within the Darktrace Threat Visualizer, programmatically via the Darktrace Threat Visualizer API, or by utilizing a model with a “tagging” action. The tags applied to a device entity, or a credential present on that device entity, are available to the Darktrace Model Engine during model evaluation, in addition to surfacing on the user interface.

A subset of tags necessary for system operation are supplied by default. Users may also define custom tags through the Darktrace Threat Visualizer User Interface or via the Darktrace Threat Visualizer API. Darktrace / NETWORK Autonomous Response adds five additional tags necessary for operation: “Antigena Compliance”, “Antigena External Threat”, “Antigena Insider Threat”, “Antigena Significant Anomaly” and “Antigena All”. Four of these tags correspond to a model category outlined above, with the exception of “Antigena All” which is equivalent to the previous four.

Darktrace / NETWORK Autonomous Response models within the four pre-defined categories contain logical criteria which will only evaluate successfully if the device entity which has met the prior criteria *is also tagged* with the corresponding Darktrace / NETWORK Autonomous Response tag for the given model category. For example, when a device is given the **Antigena Compliance** tag, it becomes eligible to receive Darktrace / NETWORK Autonomous Response responses triggered by models in the **Compliance** category.

This is referred to as *tag-based eligibility*.

The possession of a corresponding tag is therefore presumed as “eligibility” for a category of Darktrace / NETWORK Autonomous Response in default operating mode. Without eligibility, Autonomous Response inhibitors will not be effective.

Any modifications made by Customer operators to remove tag requirements from the default models outlined above are not considered within the scope of default operation. Darktrace does not recommend making modifications of this kind.

Ensuring effective tag-based eligibility across an environment is the responsibility of the Customer.

3.1.4.10. Darktrace / NETWORK Response Autonomy

Darktrace / NETWORK Autonomous Response can create actions in an “active” or “pending” state, where pending actions require human approval before any direct action is actually taken. The state actions are created in - referred to here as “autonomy” - is defined by the following factors and conditions.

3.1.4.10.1. Manually-Triggered Darktrace / NETWORK Response Actions

Manual actions are always automatically created in an “active” state.

3.1.4.10.2. Model-Triggered Darktrace / NETWORK Autonomous Response Actions

The status of actions created as a result of a Darktrace model alerts (both default Darktrace / NETWORK Autonomous Response models, and those user-created or user-modified) is defined by a combination of granular eligibility controls.

A high-level summary of the current autonomy level is provided to operators in the main Darktrace Threat Visualizer homepage. A detailed summary is also accessible from the Darktrace / NETWORK Response Quick Setup Process.

Darktrace / NETWORK Autonomous Response autonomy can be configured on a per-model basis to meet the varied requirements of each organization. This autonomy is controlled in the model configuration: models may “force human confirmation”, “force autonomous action”, or “permit autonomous action”. Actions created by models with a “force” state will be placed in a *pending* or *active* state respectively on creation, regardless of any other configuration. An action created as a result of the model with “permit autonomous action” will attempt to create an *active* action if permitted by other configuration factors but will defer to *pending* if overruled.

Templates provided as part of the Darktrace / NETWORK Response Quick Setup Process include a subset of high-severity models recommended by Darktrace to be placed in “force autonomous action” state. A configuration interface to modify these states on a per-model basis is also provided as part of the process.

Darktrace also provides a seven-day, hourly timetable (“Response Schedule”) which allows blocks of autonomous action or enforced human oversight to be scheduled. The schedule is applied across Darktrace / NETWORK, Darktrace / ENDPOINT and Darktrace / IDENTITY.

The schedule is applied to actions created by models with “permit autonomous action”. Actions created manually, or by models in a “force” state (“force human confirmation”, “force autonomous action”), are not impacted by the schedule state. Darktrace provides recommended timetables for the schedule as part of Darktrace Response Quick Setup Process templates. Access to configure the schedule directly is also provided within the Darktrace Response Quick Setup Process and from the Threat Visualizer interface.

For global organizations, the schedule can also be localized at the subnet-level. In this mode, the hours defined in the schedule are adjusted to apply in the time zone matched to the latitude and longitude information defined for the subnet in the Darktrace Threat Visualizer. Individual schedules can also be defined on a per-submaster basis in a distributed Darktrace deployment.

3.1.4.10.3. Per-Device Human Confirmation Mode

Darktrace also provides an advanced option to override autonomous actions on a per-device basis using a pre-defined tag.

3.1.4.10.4. Cyber AI Analyst-Triggered Actions

Actions created as a result of a Darktrace incident can “force human confirmation”, “force autonomous action”, or “permit autonomous action”. This setting is set on the Darktrace System Config page and applies to all actions created by Darktrace Cyber AI Analyst.

Actions created when this setting is in a “force” state will be placed in a *pending* or *active* state respectively on creation, regardless of any other configuration. An action when set “permit autonomous action” will attempt to create an *active* action if permitted by other configuration factors but will defer to *pending* if overruled.

3.1.5.Interface – Darktrace Threat Visualizer

Read-only access to view Darktrace / NETWORK Autonomous Response actions is available to all users with access to the Darktrace Threat Visualizer interface. Modification of actions within these interfaces - including the creation of new actions - is permission restricted.

3.1.5.1. Overview

The main Threat Visualizer homepage displays a high-level overview of how autonomously Darktrace Autonomous Response can act at the current moment.

This overview summarizes across all possible factors, configurations, and settings to give a true representation of how autonomously Darktrace can act at each point of the day. When collapsed, the element will indicate the overall autonomy level and how long this state is applicable for.

A high-level, exportable summary of Darktrace Autonomous Response actions over the last seven days is also available from the Threat Visualizer Response Actions Summary view. This overview provides high level statistics regarding Darktrace / NETWORK Autonomous Response activity including the action types created, human approval granted, and the meantime to activation for pending actions.

Darktrace / NETWORK Autonomous Response actions are also included in reporting outputs (Executive Threat Report, Operations Report, Cyber AI Insights Report).

3.1.5.2. Actions

Darktrace / NETWORK Autonomous Response actions are recorded in the Response Actions window within the Threat Visualizer. Actions are sorted by type, then by state. A per-device filtered view is also offered. Users with appropriate permissions may alter the state of an action from this view (clear, activate, extend, or reactivate). Darktrace / NETWORK Autonomous Response actions are also surfaced in, and can be modified from, the Darktrace Mobile App.

Pending Darktrace / NETWORK Autonomous Response actions produce a notification above the threat tray in the Threat Visualizer. The count of pending actions is also displayed in the homepage summary. Darktrace offers alerting when Darktrace / NETWORK response actions are created or change state; notifications are surfaced in the Threat Visualizer interface, in the Darktrace Mobile App, and sent via compatible alert outputs.

Devices currently undergoing actions are visually distinguished (at present, with a green highlight) across the Darktrace Threat Visualizer interface for the duration of the action. Devices with pending actions are also separately distinguished (at present, with an orange highlight) to indicate human approval is pending for a Darktrace / NETWORK response action.

Connections prevented by Darktrace / NETWORK Autonomous Response are indicated in device event logs throughout the interface. Where network traffic ingested by Darktrace is found to include Darktrace / NETWORK Autonomous Response TCP RST actions, Darktrace Deep Packet Inspection will distinguish these packets from standard TCP RSP packets with a special code ("G/g") in the Darktrace Advanced Search interface entry.

3.1.5.3. Autonomous Response Models

Darktrace / NETWORK Autonomous Response models in default operation mode will create standard model alerts in the Darktrace Threat Visualizer interface. Where an action is created as a result of a model, a pivot point from the model alert is also provided.

3.1.6. Reporting

The Darktrace Threat Visualizer "Audit Log" records changes made by operators to Darktrace Autonomous Response actions; the audit log can be exported via Syslog for extended retention.

User interactions with Darktrace / NETWORK Autonomous Response actions which alter the state (e.g. "activate", "extend", "clear") are also recorded in the action history, accessible from the Darktrace Response Actions window. Users can also be compelled to provide a free-text justification when an action state is altered, which is also displayed where applicable.

Individual action-level information is surfaced primarily in the Darktrace Threat Visualizer Response Actions window and Darktrace Mobile App.

Darktrace provides a REST API for automated retrieval of a subset of data output.

3.2. Deployment Architecture

The appropriate combination of deployment components will vary depending on the network environment, for example, the balance of physical and virtualized infrastructure, or the physical location of covered components. The appropriate deployment scenario may also change during the lifetime of the Darktrace deployment if significant changes are made the network architecture. It is the responsibility of the Customer to ensure that Darktrace maintains both continued visibility and continued *reachability* over network traffic to facilitate operation of Autonomous Response.

Darktrace strongly recommends that deployments are architected in line with best operating practice. In the event of re-architecture of the environment, Darktrace recommends Customer reaches out to the Darktrace support team to determine the impact on the operation of Autonomous Response.

3.2.1. Topology Components

Darktrace / NETWORK Autonomous Response TCP RST actions may be issued from any compatible Darktrace ActiveAI Security Platform instance. All network architecture components will attempt to send spoofed Darktrace / NETWORK Autonomous Response TCP RST packets in response to targeted connectivity.

Compatible topology components include physical Darktrace probes, physical masters, physical subordinate masters, physical Unified View instances, vSensors with appropriate configuration and Darktrace osSensors. Darktrace-hosted virtual cloud masters are located within the Darktrace cloud environment and therefore cannot issue TCP RST packets to targeted endpoints. To take action in this scenario, a firewall integration or other Darktrace topology component (e.g. as a physical or virtual probe

located within reach of the target entity) is required. For more information on the Topology Components, please refer to the Real-Time Detection section of the Product Specification.

3.2.2. Packet Traversal & Stateful Boundaries

In the majority of operating scenarios, Darktrace / NETWORK Autonomous Response TCP RST packets will be unable to cross stateful boundaries or may be limited in traversal by network routing components. This may be due to a number of factors including configuration restrictions in place within the networking component, anti-spoofing measures present in the component, or incomplete information in the network traffic sent to Darktrace which prevents it from appropriately routing the spoofed packet back to the target entity. The deployment must therefore be architected to ensure Darktrace / NETWORK Autonomous Response packets are able to traverse to all locations, including those behind stateful boundaries.

Darktrace provides mechanisms to ensure that packets can be routed to all areas of the network including dedicated “firing” network interfaces and custom-defined routing rules for spoofed packets.

Best practice recommendations are available for the implementation of Darktrace / NETWORK Autonomous Response which provide guidance on these considerations.

3.2.3. Firewall Integration

Darktrace / NETWORK Autonomous Response integrates with a selection of firewall vendors to allow Darktrace / NETWORK Autonomous Response's targeted network blocking actions to be performed by third-party firewalls. These integrations can be configured and used independently of Darktrace / NETWORK Autonomous Response's TCP Reset capabilities and are included in the Darktrace / NETWORK product.

Utilizing a Darktrace / NETWORK Autonomous Response firewall integration offers an alternative route to neutralize connections crossing the network boundary, where UDP traffic is frequently seen, or in scenarios where Darktrace / NETWORK Autonomous Response RST packets may not be able to traverse the network from the Darktrace instance to their intended destination. As such, Darktrace recommends integrating with all compatible firewalls to supplement the native TCP RST blocking to provide redundancy in case of unforeseen network changes.

3.3. Implementation

3.3.1. Licensing

Darktrace / NETWORK Autonomous Response will only take autonomous actions when a valid License Key is configured on the Threat Visualizer's System Config page. In Unified View environments, the License Key should propagate from the UV master instance to the subordinate Master instances unless specifically configured otherwise. Adding the License Key will also activate Darktrace / NETWORK Autonomous Response capabilities on all connected virtual sensors (vSensors). Instructions on adding a License Key can be found in the relevant documents in the Customer Portal.

3.3.2. Tagging & Scope

There are multiple approaches that can be taken to optimizing Darktrace / NETWORK Autonomous Responses for a given network environment. The exact methodology for deploying Darktrace / NETWORK Autonomous Response will depend both on the composition of the Customer's network and the Customer's own preference.

Darktrace / NETWORK Autonomous Response models indicate how, when and which response actions can be taken against devices. For a device to be in scope of Darktrace / NETWORK Autonomous Response it must have an "Antigena" tag applied to it. There are 5 default tags, as described above, that can be applied to devices to include them in the scope. Custom tags can be defined by modifying existing models or creating new models. Some tags may be included in models under other folders/categories, so it is not advised to modify legacy models. If the Customer decides to modify the default tags or create custom ones, Darktrace bears no liability for the consequential impact on the operation of the Product.

Additionally, Customer can also add devices under scope by using the Darktrace / NETWORK Response Quick Setup. The Darktrace / NETWORK Autonomous Response Quick Setup Process offers one click setup option to select which devices are eligible, define any activity-based overrides, and to set the timed schedule.

It is the Customer's responsibility to determine the scope of Darktrace / NETWORK Autonomous Response by the application of tags.

Darktrace recommends all devices are in scope of Darktrace / NETWORK Autonomous Response as any untagged devices will not be protected during a compromise and any activities from those devices cannot be contained.

3.3.3. Testing & Reachability

Darktrace / NETWORK Autonomous Response's ability to access different areas of the network environment ("reachability") must be tested. If reset packets cannot traverse the network from the Darktrace instance to their intended destination, customer network will not be fully covered by Autonomous Response; Darktrace / NETWORK Autonomous Response will not be able to quarantine a device or perform the full range of autonomous actions.

Testing of actions and reachability should be done for every part of the customer's network, ensuring that each separated segment (whether logically/physically etc.) is reachable by the spoofed TCP RSTs. The Customer is responsible for running these tests and implementing any remediations. Given the complexities of networks, their propensity to change over time, and the reliance on other 3rd party software/hardware/firmware, the Customer should confirm reachability across all network segments every 6 months or if there is a significant network change.

Testing can be done manually or automatically using the Darktrace / NETWORK Autonomous Response Quick Setup Process "Spot Testing" capability (recommended). Instructions on how to perform testing can be found in the relevant document in the Customer Portal.

3.3.4. Dedicated Firing Interfaces (DFIs)

Darktrace / NETWORK Autonomous Response fires reset packets from the administrative interface of Darktrace instances by default. If the administrative interface is placed in a restricted portion of the network, the reset packets may not reach their destination. To ensure that packets are able to reach destination in separated parts of the network, it is possible to configure an additional or alternative firing interface from the Console called Dedicated Firing Interface (DFI).

For every appliance, at least one DFI is recommended to be installed. Customers using non-flat networks should configure a Custom Route to send RST packets through this DFI on each appliance, using the PACKET method initially. Customers using VLANs should additionally configure the Custom Route to use VLAN tags and ensure that all expected VLAN IDs are included in the mirror feed into the relevant Darktrace appliances/sensors.

In highly segregated networks it may be necessary to use multiple DFIs in order to achieve 100% reachability. The Customer may have to physically/logically plug a DFI from multiple Darktrace appliances

across the network to bypass certain switches/firewalls or other network limitations. For vSensors, it is not usually needed to install a physical dedicated firing interface as these are mostly virtual. Additional virtual interfaces can be assigned to the vSensor VM and custom routes sent accordingly through this interface. However, it is also possible to attach a physical interface, providing the VM host machine has a spare NIC interface. This interface can then be used as a normal physical DFI plugged into a Darktrace appliance. Further details on how to configure DFIs can be found in the relevant documentation on the Customer Portal.

3.3.5. Schedule/Going Autonomous

Whether Darktrace / NETWORK can take action autonomously, or must wait for human approval, is defined by a schedule or model setting. The seven-day, hourly timetable allows blocks of autonomous action or enforced human oversight to be scheduled. For some high severity activity types, it is recommended to ensure Darktrace / NETWORK Autonomous Response can always take an autonomous action regardless of the scheduled state; in this case, an override can be configured in the model setting.

When using the Darktrace / NETWORK Autonomous Response Quick Setup, One Click Setup options are graded from least to most autonomous. In the most autonomous mode, Darktrace / NETWORK Autonomous Response will take action autonomously in all cases, requiring no human oversight. In partially autonomous operating mode, Darktrace / NETWORK Autonomous Response is restricted to taking action autonomously only outside of business hours, when there may be no one to approve pending actions. Finally, human confirmation mode compels Darktrace / NETWORK Autonomous Response to request human confirmation before acting in all but the most severe cases.

But for models that block trigger on privileged user activities, Darktrace recommends that Darktrace / NETWORK Autonomous Response models to be set to autonomous every hour of every day (24/7). Human Confirmation mode should be considered for initial setup as it allows customizations to be made based on the behavior Darktrace is seeing in different networks but is not the designed end-state for Darktrace / NETWORK Autonomous Response. Please see above for the possible methods of implementation that allow gradual implementation, via a hybrid deployment, to reach a fully autonomous deployment. A hybrid deployment is considered any that is neither fully autonomous nor fully human confirmation mode.

The Customer is responsible for configuring Darktrace / NETWORK Autonomous Response to run fully autonomously. Fully autonomous mode is the end goal of all Darktrace / NETWORK Autonomous Response deployments; reaching a fully autonomous state where the AI can take action whenever unusual or concerning behavior is detected, without the need for human oversight. This mode lends itself to a minimal-interaction workflow, where Customer may infrequently modify actions through the Darktrace Threat Visualizer interface, API, or Darktrace Mobile App, but on the whole leave Darktrace / NETWORK Autonomous Response to operate with little intervention.

3.3.6. External Alerting & Mobile App

Darktrace supports a number of Workflow Integrations directly or through generic formats such as syslog; all integrations can be reviewed on the Modules page of the Darktrace System Config page. Darktrace recommends external alerting from the UI to be configured for Darktrace / NETWORK Autonomous Response actions.

Darktrace / NETWORK Autonomous Response action alerts are a new alert format for Darktrace 6.1, providing even greater visibility over Darktrace Autonomous Response activity and notifying users when the AI is waiting for Customer approval. Each alert includes a link back to the relevant action to quickly grant approval or alter the action in some way. By default, alerts are created when an action is pending and Darktrace needs human authorization to step in but can be configured for any change of state. For Darktrace environments running in a fully autonomous mode, customers who wish to keep abreast of

Darktrace activity can receive alerts when autonomous actions are created automatically, or when they are cleared or expired.

Darktrace also recommends using the mobile app and configure it for Autonomous Response notifications. The Darktrace mobile app allows users to easily access Darktrace alerts when they are on the move. This allows alerting similar to that mentioned above but also allows interaction with response actions from the app itself. This improves ease of interaction with Darktrace / NETWORK Autonomous Response actions for the Customer. Find information to configure the mobile app on the Customer Portal.

3.3.7.Active Integrations

Darktrace provides a wide range of components and integration methods that extend Darktrace / NETWORK Autonomous Response capabilities across various vendors - and ensure that Darktrace fits seamlessly into any existing security stack. Darktrace's Autonomous Response component offers a selection of direct integrations with enterprise firewall models. In scenarios where RST packets may not be able to traverse the network from the Darktrace appliance to their intended destination, or UDP traffic is frequently seen, firewall integration can ensure that connections are neutralized at the boundary. This is particularly useful in network environments with multiple internal boundaries, or where firewall operators desire more control of Darktrace / NETWORK blocking capabilities. Details on what integrations Darktrace offer can be found in the Customer Portal.

Darktrace recommends customers integrate with all available EDR and firewall third parties to supplement the native spoofed TCP RST blocking mechanisms. This allows blocking of UPD traffic but also adds redundancy for complex networks where 100 % reachability is still being addressed and to ensure no single point of failure.

Firewall integrations vary significantly between vendors and ensuring operational integration is a Customer Responsibility. In most cases, lists are required to specify the type of block the firewall will undertake on behalf of Darktrace. For such firewalls, wherever possible, Darktrace recommends configuring 3 lists in the Darktrace module and in the firewall itself. These should correspond to:

- block all outgoing connections from a source IP,
- block all incoming connections to a destination IP, and
- block matching connections (all connections from a source IP to a destination IP).

Configuring these three lists allows Darktrace / NETWORK Autonomous Response to leverage the firewall actions without significantly reducing the precision of the native spoofed TCP RSTs.

3.3.8.Model Editing

As stated, edits made to Darktrace / NETWORK Autonomous Response models are made at the Customer's own risk. Darktrace recommends that the best practice is to edit the underlying real-time detection models, if necessary, rather than the Autonomous Response models. In the event automated scripts and systems unintentionally trigger models, it is recommended that any defeating of the underlying detection is conducted as precisely as possible.

Autonomous Response is configurable to the Customer's security-availability requirements; however, the highest level of security is achieved when no model tuning is conducted. Darktrace only recommends defeating detection triggers for response models or actions that the Customer deems an impact to business continuity and at the Customer's discretion. Note that there will always be some actions against benign activity, but precise actions should not cause business impact. If business processes would be impacted by the action taken or suggested, then Darktrace recommends addressing this.

In most deployments, the data collect from one month's worth of Autonomous Response model alerts provides sufficient insight into the impact of actions and their frequency, to allow for informed decisions to be made regarding tuning. For details on the possible edits and how to implement them, see the relevant documentation on the Customer Portal.

3.4. Administration

Darktrace strongly recommends that deployments are architected in line with best operating practice. This includes administration of physical implementations or networking implementations outside the Darktrace Threat Visualizer environment, and the administration of device eligibility, action autonomy, and other components within the Threat Visualizer platform.

Darktrace provides mechanisms to assist with the administration of Darktrace / NETWORK Autonomous Response within the Darktrace Threat Visualizer. These mechanisms are outlined below.

Deployments should be tested frequently to ensure that Darktrace / NETWORK Autonomous Response retains access to control all connectivity as desired. This responsibility for regular testing and resolution of issues identified during tests lies with the Customer. Please refer to "Reachability" below.

3.4.1. Darktrace / NETWORK Response Quick Setup

Darktrace provides a "Quick Setup Process" (hereafter QSP) for Darktrace / NETWORK Autonomous Response. This process guides users through enabling Darktrace / NETWORK Autonomous Response, the selection of eligible devices, the types of activity Darktrace / NETWORK Autonomous Response is permitted to target, and the level of autonomy available to Darktrace / NETWORK response in each case. This process offers both templated operating scenarios and granular configuration steps for each deployment aspect.

The QSP is intended for standard Darktrace / NETWORK Autonomous Response deployment and operation. The QSP does not include initial roll-out processes for Darktrace / NETWORK Autonomous Response routing or integrations - these elements must be completed before attempting the process.

3.4.2. Device Eligibility

The responsibility to ensure that devices are suitably tagged - and remain tagged - lies with Customer. Devices which are not tagged for Darktrace / NETWORK Autonomous Responses will not have actions taken against them in default operating mode.

Darktrace provides a mechanism for automatic device tagging via the Darktrace / NETWORK Response Quick Setup (method A). All templated deployment scenarios provided within the Response Quick Setup include a minimum recommended tag eligibility for devices. Eligibility is divided between client and server devices. An advanced option to define eligibility rules on a *per device type* or *per subnet basis* is also offered within the Response Quick Setup.

Separate to the Response Quick Setup, Darktrace also provides a default subset of customizable models (method B) which can be modified to opt devices into Darktrace / NETWORK Autonomous Response coverage.

These mechanisms (A & B) for defining tag eligibility are rules-based and will automatically apply tags to newly observed devices which meet the criteria.

Tags can also be added manually to individual devices using any mechanism outlined above (in “Tag-Based Eligibility”). Manual tagging is not recommended as the primary eligibility method; the static nature of this eligibility is unlikely to keep step with asset or network changes over time.

Eligibility should be regularly reviewed to ensure it has remained in-step with your network environment. Darktrace provides mechanisms to review the breadth of tag eligibility coverage within the Darktrace Threat Visualizer.

3.4.3. “Reachability”

Reachability refers to the ability for Darktrace / NETWORK Autonomous Response to reach all connectivity targeted by actions.

The responsibility to ensure reachability is tested regularly and any factors impacting reachability are resolved lies with the Customer operator. Reachability testing can be performed manually, or via the mechanisms provided by Darktrace.

Darktrace / NETWORK Response Quick Setup is accompanied by an automated testing system in technical preview. This reachability *spot tester* will take quarantine actions against nominated devices in each subnet over the duration of test; results are then reported in the Threat Visualizer, and Darktrace / NETWORK Autonomous Response will also suggest possible reasons for unsuccessful tests.

Testing Darktrace / NETWORK Autonomous Response is an iterative process - every subnet should be tested to ensure that Darktrace / NETWORK Autonomous Responses can reach all eligible devices if the need arises. Network configurations change frequently; tests should also be repeated regularly to ensure routes have not been severed by new stateful boundaries.

4. Requirements

4.1. Physical Install Requirements

4.1.1. Network Access Ports

It may be necessary to configure your network security devices to permit access to and from the Darktrace physical appliances for administration and other services. The required connections below must be enabled before attempting initial setup:

CONNECTION	PORT	DIRECTION	REQUIRED?
Threat Visualizer and web configuration. Communication between master and probe/Unified View.	443 (TCP)	Inbound	Required
Console application and file transfer via SFTP ¹	22 (TCP)	Inbound	Required
Network Time Protocol	123 (UDP)	Outbound	Required
Syslog ingestion of mapping data	Various, see below	Inbound	Optional
DNS querying	53 (TCP & UDP)	Outbound	Optional
Scheduled Backups (preferred mode - SCP/SMB/S3-compatible service)	22 (TCP), 445 (TCP) or 443 (TCP)	Inbound	Optional
Call-Home	443 ² (TCP)	Outbound	Optional

Connectivity between a Darktrace Unified View instance and any subordinate Masters requires a minimum bandwidth of 5Mbps.

4.1.2.Call Home

Outbound firewall exceptions may be necessary for the appliance to connect to the Call-Home service. Each instance is provided with a unique Call-Home endpoint, which can be retrieved from the console of the instance after boot.

4.1.3.System Backups

Automatic software backups can be configured and exported to an external location over SCP, SMB or to a S3-compatible location. Please ensure the selected backup protocol/mode is allowed between the backup storage location and the Darktrace appliance. It is a Customer responsibility to configure backups.

4.1.4.Syslog Input

Inbound syslog ingestion is supported on several port/protocol combinations. The most suitable is dependent upon your network environment and the Darktrace components deployed, such as the supported export formats of the syslog sender or whether syslog must traverse untrusted networks.

CONNECTION	PORT	DIRECTION	ENCRYPTION
Syslog input to standalone master or probe	1514 (UDP or TCP)	Inbound	Unencrypted
Syslog input to standalone master or probe	6514 (TCP)	Inbound	TLS / SSL
Syslog input to Unified View	2514 (UDP or TCP)	Inbound	Unencrypted
Syslog input to Unified View	7514 (TCP)	Inbound	TLS / SSL

The default ingestion port can be altered in the host variables of the appliance after configuration. Unused syslog ports can also be disabled by the customer if desired for compliance purposes.

4.1.5. Powering the Appliance

Appliances shipped with Darktrace Threat Visualizer 4.0.7 and above will auto-boot on power. It is strongly recommended that redundant and uninterruptible power supplies are used to maintain service in the event of primary power failure.

For appliance models that contain dual-power supply units for redundancy, both supplies must be connected at all times. The units will charge when the power supplies are connected.

Some models may contain hardware RAID battery backup units (BBU's). In common with other RAID battery backup units, and to ensure data integrity, it is recommended that the units have at least 20% charge before full disk access can be assured.

4.2. Darktrace Cloud Region Availability & Security

If the customer organization is subject to regional or geographic restrictions on data flow that would prevent use any of the following cloud regions, the customer must disclose this to Darktrace. The customer must enable network access from probe components to the relevant cloud region.

For security reasons, the cloud-hosted master will not accept or ingest unencrypted data. Network traffic, log data, and other data types must be sent over a secure channel, such as a vSensor operating in an approved mode.

4.2.1. AWS Regions

Darktrace offers AWS cloud-based deployments hosted in Europe (AWS region "eu-west-1" or "eu-west-2"), the United States (AWS regions "us-west-1" and "us-west-2"), Canada (AWS region "ca-central-1"), Singapore (AWS region "ap-southeast-1") or Australia (AWS region "ap-southeast-2").

REGION	IP ADDRESS	DNS ENTRY
US (1)	52.9.179.107, 54.177.70.47	cloud-nat-usw1.darktrace.com
US (2)	54.187.177.155, 44.224.231.203, 34.215.32.12	cloud-nat-usw2.darktrace.com
Canada	15.223.16.1, 3.97.36.106, 3.98.161.226	cloud-nat-cac1.darktrace.com
EMEA (Ireland)	52.51.139.68, 54.73.200.146, 46.137.35.194	cloud-nat-euw1.darktrace.com
EMEA (UK)	18.132.236.38, 18.134.166.226, 18.169.92.72	cloud-nat-euw2.darktrace.com
APAC (Singapore)	52.220.237.248, 54.255.89.86, 54.255.12.109	cloud-nat-apse1.darktrace.com
APAC (Australia)	3.24.26.120, 52.64.37.154, 13.54.22.223	cloud-nat-apse2.darktrace.com

4.2.2. Microsoft Azure Regions

Darktrace can offer Azure cloud-based deployments hosted in EMEA (Azure regions "UK South", "West Europe", "UAE North" or "South Africa North"), the United States (Azure region "East US"), Canada (Azure region "Canada Central"), South East Asia (Azure region "Southeast Asia") or Australia (Azure region "Australia East").

REGION	IP ADDRESS	DNS ENTRY
US (East)	52.170.164.120, 20.62.136.124, 20.62.143.143	cloud-nat-eastus.darktrace.com
Canada	52.139.10.121, 20.200.75.205, 20.200.74.241	cloud-nat-canadacentral.darktrace.com
EMEA (UK)	20.49.143.39, 20.77.146.47, 20.77.145.243	cloud-nat-uksouth.darktrace.com
EMEA (EU)	20.61.9.184, 20.86.224.59, 20.86.224.40	cloud-nat-westeuropa.darktrace.com
APAC (Australia)	20.193.44.157, 20.53.93.129, 20.53.93.131	cloud-nat-australiaeast.darktrace.com
APAC (SEA)	20.197.98.133, 20.43.153.225, 20.43.153.237	cloud-nat-southeastasia.darktrace.com
UAE North	20.233.136.200, 20.233.136.204, 20.233.136.214	cloud-nat-uaenorth.darktrace.com
South Africa North	20.87.94.11, 20.87.94.70, 20.87.94.88	cloud-nat-southafricanorth.darktrace.com

4.2.3. Darktrace Probe Requirements

4.2.3.1. vSensor requirements

The vSensor ingests and processes network traffic before sending it to the Darktrace Master instance. The size of data transferred across the network is approximately 1-4% of the traffic ingested by the vSensor. The table below shows the required network ports for communication.

PURPOSE	PORT / PROTOCOL	DIRECTION	REQUIRED?
Contact packages.darktrace.com for updates	443/TCP (HTTPS)	Outbound	Required
Contact packages-cdn.darktrace.com for updates	443/TCP (HTTPS)	Outbound	Required
Communication with associated osSensors	443/TCP (HTTPS) & 80/TCP (HTTP)	Inbound	Required (when osSensors deployed)
Remote management of Cloud Deployments	22/TCP (SSH)	Inbound	Required

Ubuntu 20.04 Focal Server is required for Darktrace vSensor 5.1 and above. This operating system is supported by Linux. Darktrace will align to update the vSensor to newer operating system before end-of-life support by Linux.

vSensors require at least two CPU cores but works best with three or more. 2GB RAM is required, however 3.75-4 GB RAM is a highly recommended minimum for substantial performance.

Deep Packet Inspection workers scale automatically, it is recommended to keep at least a 2GB per core ratio, with 4GB per core being optimal.

PCAPs are stored on the vSensor disk. A minimum disk of 20GB is recommended, retention is dependent on available storage and traffic ingestion, so better retention can be achieved with more disk space provisioned.

4.2.3.2. osSensor Requirements

To perform the installation process for a Darktrace osSensor, there must be a Darktrace Master appliance running the most recent version of the Darktrace Threat Visualizer, and the IP of a configured vSensor must be able to communicate with the intended osSensor location.

4.2.4. Darktrace Server Agent Requirements

The Server Agent must be installed on each instance of the Server hosting centralized applications. The TSA customer server on which the Agent is deployed location(s) must be able to contact the Darktrace Master instance in order to send connection information.

5. Customer Responsibilities

Notwithstanding any additional responsibilities set out in any other section of this Product Specification, and without limitation, Customer's responsibilities are to ensure that:

- Darktrace / NETWORK Real-Time Detection has been deployed in line with Darktrace best practice recommendations and, if changes are made to the network configuration, ensure the deployment continues to meet best practice recommendations across its lifecycle.
- Darktrace maintains visibility over all network traffic to be monitored.
- The network traffic provided to Darktrace is of suitable quality and does not contain duplication, fragmentary data, or is incomplete in any fashion.
- The network traffic provided to Darktrace components is within the relevant specification of those components and that deployment components are not overloaded.
- System health issues are monitored using the built-in features of the Darktrace / NETWORK product, and if alerts arise, are addressed and rectified in a timely manner.
- The subnets detected by Darktrace / NETWORK Real-Time Detection are configured to be tracked by the optimum available identifier.
- End-user access is managed securely and within recommended best practice.
- Changes made to components or overall system configuration by users of the Darktrace platform do not impact the system's ability to function, or otherwise degrade service.
- Darktrace / NETWORK Autonomous Response has been deployed in line with Darktrace best practice recommendations and, if changes are made to the network configuration, ensure the deployment continues to meet best practice recommendations across its lifecycle.
- Darktrace / NETWORK Autonomous Response has - and maintains - reachability to all devices intended for Darktrace / NETWORK Autonomous Response actions. Where reachability is impeded, it falls within the customer's responsibility to identify the impedance and ensure that reachability is restored.
- Where Darktrace / NETWORK Autonomous Response TCP RST packets are unable to cross stateful boundaries or are limited in traversal by network routing components, to ensure that a dedicated firing interface or alternative configuration is in place to route packets to all network locations.

- There is knowledge and understanding of customer network topology, systems, etc. such that appropriate best practices and necessary modifications to Autonomous Response can be implemented.
- Darktrace / NETWORK Autonomous Response is enabled in the appliance console (where relevant), in the deployment settings, and in any other applicable locations to ensure actions can be taken.
- Darktrace / NETWORK Autonomous Response is granted necessary autonomy to enact actions, or, where human confirmation mode is partially or fully in place, ensure that actions are activated by human operators in a timely manner.
- Suitable workflows and processes are in place to remediate any cyber-attacks/compromises that are contained by Darktrace / NETWORK Autonomous Response
- Devices are opted appropriately into Darktrace / NETWORK Autonomous Response using tags.
- Alterations are not made to Darktrace / NETWORK Autonomous Response models, or other Darktrace models utilized by Darktrace / NETWORK Autonomous Response, which would significantly impede the ability to create targeted actions.
- Darktrace / NETWORK Autonomous Response actions are not created excessively by custom models or existing models modified to contain actions.
- All desired integrations with third party firewalls and EDRs are configured to enable action by Darktrace on all desired devices.
- The outputs of the Darktrace / NETWORK Autonomous Response are monitored, and its use complies with all applicable laws and regulations.

Examples of configuration changes which may result in degraded service include the exclusion of traffic or devices from deployment scope, the incorrect modification of subnet tracking configuration, misidentification of device types, the creation of excessively overactive models, and the modification of existing models to the extent that alerting criteria can no longer be met.

The RACI matrix below details the accountability and responsibility between Darktrace and Customer where no Additional Services have been purchased. Customer should be familiar with their existing security & network infrastructure and must be able to perform the actions outlined.

Role	Activity/Responsibility	
R – Responsible	The party is responsible for implementation of the activity; owns the problem or project	
A - Accountable	Right to make decisions, signs, and improves work	
C – Consulted	Has information, resources, and/or capacities necessary to assist the job	
I – Informed	Must be informed of the results, but does not need to be consulted	
Activity/Responsibility	Darktrace	Customer
Implementation, Administration & Maintenance		
Product Implementation, Administration and Maintenance	CI	RA

Provision of guides to support Implementation, Administration & Maintenance	RA	CI
Product Licensing	RA	CI
Maintain coverage, including visibility, tagging and eligibility within product license	CI	RA
Monitor and resolve system status alerting	CI	RA
Administrate access to the Threat Visualizer	CI	RA
Administrate access to the Darktrace Customer Portal	CI	RA
Performing regular testing to ensure continued reachability	CI	RA
Managing active integrations and model editing	CI	RA
Ongoing maintenance of Autonomous Response configurations	CI	RA
Physical components		
Provide shipping and appliance configuration information	CI	RA
Provide appliances, subject to customer-provided information	RA	CI
Configure network settings for appliance	CI	RA
Maintain & modify traffic feeds as required, within appliance specifications	CI	RA
Modify Firewall rules to allow communication with Darktrace infrastructure (Call Home)	CI	RA
Modify firewall rules to allow communication between Darktrace Master and Probe components	CI	RA
Virtual components		
Make software available as detailed in customer contract terms	RA	CI
Download software from Darktrace infrastructure	CI	RA
Ability to access and deploy new virtual machines to support sensor components	CI	RA
Modify network access rules to allow communication between Darktrace sensors and Master instance	CI	RA

Modify network access rules on endpoint to allow communication & updates from Darktrace infrastructure	CI	RA
Maintain & modify traffic feeds as required	CI	RA

6. Considerations

Where source data is low quality, represents only a subset of all network activity, is delayed by external factors, or otherwise is incomplete, the quality of Darktrace analysis will be impacted. It is therefore imperative that data ingestion is configured comprehensively and to a high standard by the client. Analysis may also be impacted by overloading, such as that observed when traffic throughput significantly increases beyond scoped levels.

Each deployment has a specified capacity, determined by reference to the Usage Metrics band that has been purchased. Exact performance may vary depending on several factors that may be unique to the network being analyzed. Staying within Usage Metrics band allows for optimal Darktrace performance.

In the event that Customer has provided incorrect or incomplete sizing information or goes allows their usage to exceed the Usage Metrics band, Darktrace bears no liability should the actual traffic in Customer environment exceed the sizing information and degrade service.

In the event that actual traffic in Customer environment exceeds the Fencing Metrics as calculated for the relevant Usage Metrics band, Customer acknowledges this may result in, without limitation: overloading, inconsistent service, delayed Autonomous Response actions, a processing queue for new traffic, packet drops, and unsupervised learning being turned off.

Environments which have not been adequately configured to track devices by the most appropriate, consistent factor, or where tracking has not been updated in line with network changes will be unable to develop long-term 'pattern of life' behavioral profiles for unusual activity detection.

Where a Darktrace environment undergoes significant system load - typically due exceeding recommended operational limits - "High performance mode" may be activated. This mode disables a subset of high load 'pattern of life' classifiers to reduce CPU/memory usage within acceptable limits.

Darktrace / NETWORK Autonomous Response can only take automatic actions when granted appropriate autonomy. If actions are left in a pending state, Darktrace / NETWORK Autonomous Response cannot perform the required actions to mitigate connectivity.

Darktrace / NETWORK Autonomous Response TCP RST packets are typically unable to cross stateful boundaries without configuration specific to customer environment and may also be limited in traversal by network routing components. Darktrace provides best practice guidance to mitigate these limitations. The customer must ensure that these guidelines are followed to provide routing of response packets to all locations.

Where the source and destination device are located in close physical proximity, the round-trip time for the ingestion of the mirrored traffic into the Darktrace instance and subsequent RST packet creation may prevent successful connection interruption.

Darktrace / NETWORK Autonomous Response will be unable to target short-lived connectivity which has completed before traffic is received by the Darktrace instance. Where Darktrace / NETWORK Autonomous Response creates an active action, but no subsequent connectivity which meets the criteria is observed during the duration, no RST packets will be issued despite the action creation.

Darktrace / NETWORK Autonomous Response TCP RST actions are not compatible with UDP traffic.

In the event that a bug is identified, the Customer must notify Darktrace in writing. Darktrace agrees to make commercially reasonable efforts to fix or provide a workaround for any critical bugs but makes no guarantees regarding the time required to resolve non-critical issues.

7. Roles

Customer Role	Responsibility
Darktrace Product Owner / Project Manager	Coordinates Customer resources as necessary. Serves as the primary point of contact between Customer and Darktrace. Drives communication from the Customer side. Serves as the point of escalation for issue resolution and service-related issues.
Users of Darktrace Platform	Leverages Product functionality; monitors & actions output. Provides feedback to other Customer & Darktrace roles
Network/System Administrator	Ensures Darktrace / NETWORK coverage is maintained. Performs regular checks of data inputs. Responds accordingly to System Status alerts to ensure optimal traffic quality and delivery. Maintains call-home connectivity with Darktrace (where enabled).
Customer Portal Primary User	Manages customer portal access and contact information for all other customer roles. Ensures that Service Contacts are verified.

Darktrace Role	Responsibility
Customer Success Manager	Oversee customer's experience using Darktrace. Act as customers' sponsor & escalation path internally within Darktrace.
Darktrace Customer Support	Provides Support Services as per Master Customer Agreement
Account Executive	Facilitate commercial arrangements between Customer and Darktrace
Solutions Engineer	Provides technical expertise to support Account Executive and Customer Success Manager, regarding expansions or changes of coverage.

Darktrace / OT

- 1. Introduction 41
- 2. Product Features 41
 - 2.1. High Level Summary 41
 - 2.2. Data Retrieval and Ingestion 42
 - 2.3. Darktrace Analysis 42
 - 2.4. Output 44
 - 2.5. Interface 45
 - 2.6. Reporting 45
- 3. Deployment Architecture 45
 - 3.1. Darktrace OT Architecture scope 45
- 4. Implementation 46
- 5. Administration 46
 - 5.1. Updates to Threat Intelligence Datasets 46
- 6. Requirements 46
- 7. Customer Responsibilities 48
- 8. Considerations 49
- 9. Roles 50

1. Introduction

Darktrace / OT was previously named DETECT/OT, DETECT & RESPOND/OT, and Industrial Immune System. This specification applies for Customers who have purchased the product under any of the above naming conventions.

2. Product Features

2.1. High Level Summary

Darktrace / OT evaluates activity across both Information Technology (IT) and Operational Technology (OT) environments to identify behavior that deviates from a baseline “normal” state. This baseline is created through continual analysis of network connections and other supplied data events. The system profiles individual devices based upon their activity, and the activity of those it deems “peers” due to similar behavioral activity (“modelling”). Alerts are raised when specific criteria or a minimum threshold of unusual behavior is met. The output is surfaced to operators in the Darktrace Threat Visualizer interface and other relevant sub-pages for investigation and resolution.

2.2. Data Retrieval and Ingestion

All material outlined in sections “2.1.2 Data Retrieval and Ingestion” and “3.1.2 Data Retrieval and Ingestion” of the Darktrace / NETWORK specification - including sub-sections - is applicable here.

2.3. Darktrace Analysis

All material outlined in section “2.1.3 Darktrace Analysis” and “3.1.3 Darktrace Analysis” of the Darktrace / NETWORK specification - including sub-sections - is applicable here.

2.3.1. Darktrace Deep Packet Inspection

In addition to those protocols supported by Darktrace / NETWORK, Darktrace / OT supports a range of protocols specific to Operational Technology environments. A list of currently supported OT protocols is available on request.

2.3.2. Darktrace ‘Pattern of Life’ Analysis and Classification

Classification may be tuned using metrics or weightings relevant to specific OT device types or protocol behaviors.

2.3.2.1. ICS Devices and Device Types

During analysis, Darktrace / OT may classify observed devices as “ICS devices”. This designation is based on communication patterns and device behavior. Exact criteria are proprietary and subject to change.

Devices designated as “ICS devices” are eligible for additional analysis, for example, risk and weakness analysis by Darktrace OT Risk Management.

This designation is not limited to devices assigned a Darktrace / OT device type.

Darktrace / OT also supports limited modelling and analysis for certain compatible non-IP devices (for example, serial devices). These devices are distinguished using prefixes that indicate the gateway or associated device through which their communication is observed. Non-IP devices are not enumerated in asset counts.

2.3.3. Darktrace Real-Time Detection Models Framework

Darktrace / OT environments receive both the default Darktrace / NETWORK model library and an additional model set specific to OT use cases. These OT models address behaviors and threats relevant to industrial environments and are mapped to the MITRE ATT&CK ICS framework. Updates to these models follow the same process as updates to other default models.

2.3.4. Operational Alerts

In addition to standard model alerts, Darktrace / OT evaluates a distinct category of alerts known as Operational Alerts. Operational Alerts highlight operational issues that do not represent cyber threats. Operational Alerts operate within the existing model framework and can be surfaced through existing alert mechanisms.

Operational Alerts are not subject to Cyber AI Analyst analysis.

2.3.5. Cyber AI Analyst

Darktrace Cyber AI Analyst performs a meta-analysis upon the previous layers of analysis described. Please refer to “Cyber AI Analyst” in the Product Agnostic section below.

2.3.6. Additional Integrations

Darktrace / OT offers integrations and analysis modules tailored specifically for OT environments. These include integrations with asset management systems, contextual threat intelligence sources (such as passive CVE tracking and end-of-life information), and active asset identification (Active ID).

These integrations are configured on the Darktrace Threat Visualizer “System Config” page.

2.3.6.1. Active Device Identification

Active Device Identification (Active ID) enables users to directly query devices tagged and modelled as an ICS device type, using a supported OT protocol to collect any available asset information. The available protocols include: BACNet, CIP, DNP3, MMS, Modbus, OPC UA, S7, SNMP and HTTP.

A “Routing” menu is made available for users to specify which Darktrace sensor and interface should be used to initiate a connection. Unused Darktrace analysis interfaces can be reconfigured for this purpose.

Active Device Identification can be configured in various modes, based on the scope of assets and their anticipated sensitivity to active querying. All modes support one-off queries and a scheduled scan option.

2.3.6.2. Single Device Mode

The “Single Device” mode allows a device to be targeted by the IP address specified in the Address field. The Darktrace instance and interface used to target the device can also be defined within the “Routing” option. Within the Single Device mode menu, a protocol can be selected as well as the option to override the default port associated with the protocol, before initiating a connection attempt. For the BACNet protocol, an instance number for the protocol can be inputted, whilst for the SNMP protocol, and SNMP Community can be inputted. The default SNMP community is “public”.

2.3.6.3. Industrial Device Mode

The “Industrial Device” mode will establish a connection between a Darktrace Sensor and all OT devices within scope and send an ICS Discover request. OT devices are devices observed communicating on the network using an OT protocol. Each device will be targeted using the OT protocol or protocols it has been observed using in the last two weeks. The support protocols have been listed in section 3.1.1 “Active Device Identification”.

The Industrial Devices mode allows specific devices to be included or excluded based on device tags. Devices can also be excluded based on subnets (using a comma separated list for multiple subnets).

2.3.6.4. SNMP mode

Similar to the previous options, SNMP mode allows users to configure and SNMP scan with user defined parameters to include or exclude devices based on tags or subnet ranges

2.3.6.5. HTTP mode

The HTTP mode allows a single device to be identified in a single specific HTTP URL.

2.3.6.6. Vulnerability Tracker

The Darktrace ICS Vulnerability Tracker identifies common vulnerability and exposures (CVEs) present of any OT device observed by Darktrace/OT. The tracker can operate in an online or offline mode. Online requires an active connection to a vulnerability database hosted by NIST, whilst the offline option queries a database hosted and maintained by Darktrace on the Master appliance.

2.3.6.7. End of Life Tracker

The Darktrace end of life tracker monitors the lifecycle status of compatible devices and operating systems. The tracker scans for devices in both OT and IT environments and labels devices with a corresponding End of Life tag and then feeds this information into the OT Risk Management module to provide greater context for more accurate risk scoring.

Supported OT Vendors as of Darktrace 7.0 are: Siemens, Schneider, Rockwell, Mitsubishi and Omron. Supported IT Operating Systems are: Windows, MacOS and Linux (Ubuntu, Red Hat and Debian only).

2.4. Output

All material outlined under sections “2.1.4 Output” and “3.1.4 Output” of the Darktrace / NETWORK specification - including sub-sections - is applicable here.

2.4.1. Operational Alerts

Darktrace / OT generates a third category of alert known as Operational Alerts, which identify OT operational issues that are not cyber threats. Operational Alerts can be delivered through existing alerting mechanisms or viewed in the Operational Overview interface.

2.4.2. Darktrace / OT Risk Management + CVEs and Mitigations

Users of OT Risk Management may create external alerts to track patching progress for detected CVEs or to monitor the implementation of recommended mitigations. Alerts are created on demand from the OT Risk Management interface and are not sent automatically.

Mitigations are prioritized, actionable recommendations designed to reduce risk, such as closing attack paths or hardening configurations. Each mitigation includes a description, MITRE ATT&CK mapping, impacted assets, and options for export or ticket creation.

CVEs (in the context of OT Risk Management) are detections of vulnerabilities that originate from Darktrace analysis, external scanners, or manual uploads. Darktrace enriches CVEs with context such as affected devices, CVSS/EPSS scores, the Darktrace Vulnerability Score, and potential business impact. Users may accept associated risks or raise external tickets for remediation.

Mitigations or CVEs marked as resolved in the external platform can sync back to OT Risk Management, and vice-versa.

2.5. Interface

All material outlined in sections “2.1.5 Interface” and “3.1.5 Interface” of the Darktrace / NETWORK specification - including sub-sections - is applicable here.

Darktrace/OT uses the Threat Visualizer as the primary investigation interface. OT-specific information including device behavior, protocols, events, and models are embedded within the existing interface. OT-specific modules and integrations are enabled and managed in System Config.

Operational Alerts are presented separately within the Operational Overview. Both the Operational Overview and OT Risk Management are provided as dedicated interfaces (see below).

2.5.1.Operational Overview

The Operational Overview consolidates OT asset information, operational alerts, and key risk indicators into one view. It provides access to Operational Alert investigation and presents asset statistics, discovery outputs, and (where licensed) risk summaries derived from OT Risk Management or Proactive Exposure Management.

The interface serves as a central location for reviewing asset inventories, operational health indicators, and risk summaries.

2.5.2.OT Risk Management

The OT Risk Management interface establishes a baseline of OT risk, maps likely attack paths, highlights high-impact entities and vulnerabilities, and prioritizes mitigations mapped to MITRE ATT&CK frameworks. Insights are contextualized using risk factors, potential business impact, and potential interest from APT groups.

The interface provides advanced searching across identities, accounts, devices, and vulnerabilities, as well as reporting, ticketing, and configuration options.

2.6. Reporting

All material outlined in sections “2.1.6 Reporting” and “3.1.6 Reporting” of the Darktrace / NETWORK specification - including sub-sections - is applicable here.

Device information may also be visualized within the Operational Overview, as described in section 1.5.1.

3. Deployment Architecture

All material outlined in sections “2.2 Deployment Architecture” and “3.2 Deployment Architecture” of the Darktrace / NETWORK specification - including sub-sections - is applicable here.

In addition to standard physical probes, a ruggedized, specialist hardware probe - the “IMP” - is available for deployment in OT environments.

3.1. Darktrace OT Architecture scope

Darktrace / OT's Master-Probe deployment topology facilitates a wide scope of traffic collection, reaching down to basic controllers (level 1 of the Purdue Model). The separation of roles of each interface of Darktrace's sensors ensures network segmentation and separation of IT and OT networks remain intact.

By default, analysis interfaces, though physically connected to the OT network switches, are not configured with an IP allowing them to communicate with assets within the network. Additionally, analysis interfaces have no means of accessing the management interface and/or management console, meaning if the management interface was configured in a different network to where data was being captured, there would be no logical bridge between the two environments.

4. Implementation

All material outlined in section “2.3 Implementation” of the Darktrace / Network specification – including sub-sections – is applicable here.

The Darktrace / OT enrichment modules are incorporated into the master and probe sensors and require no additional dedicated installations. Configuration of all modules is managed within the system configuration menu, accessible through the Threat Visualizer UI.

5. Administration

All material outlined in sections “2.4 Administration” and “3.4 Administration” of the Darktrace / NETWORK specification - including sub-sections - is applicable here.

5.1. Updates to Threat Intelligence Datasets

Darktrace / OT integrations that rely on external datasets (for example, the NIST National Vulnerability Database) attempt automatic updates where possible. If internet connectivity is unavailable or direct updates are not configured, Darktrace will supply updates through existing offline channels.

Further detail is provided in the product-agnostic section on “Operational Change Management and Version Updates”.

6. Requirements

- All Darktrace instances intended to analyze OT traffic must have OT-specific features enabled, using conversion codes available on request from Darktrace.
- Darktrace / OT integrations must be configured according to their respective documentation (available on the Darktrace Customer Portal), with required credentials, API keys, service accounts, or connectivity provided as needed. These requirements include, but are not limited to:
 - For operation of the PAS Cyber Integrity integration:
 - The PAS platform must be reachable by the Darktrace instance and must provide asset data in the expected format.
 - PAS device records must include accurate identifiers (for example, MAC and/or IP) to allow correlation with Darktrace / OT devices.
 - For operation of the IEC 62443-3-3 Compliance Module:
 - The module is enabled in System Config and the Darktrace environment must contain sufficient historical device and alert data for evaluation.
 - Routine runs of the compliance check are initiated as needed to keep results current.

- For operation of the Vulnerability Database Integration:
 - Outbound connectivity (and proxy access if required) must be available to retrieve vulnerability data.
 - Update schedules must be configured and permitted to run to maintain current vulnerability information.
- For operation of the ICS Vulnerability Tracker:
 - Devices must expose or provide the details required for CVE matching (for example, model or firmware version).
 - The tracker must be enabled and must receive updated CVE data from the Vulnerability Database Integration.
- For operation of the Active Device Identification module:
 - The Darktrace appliance and selected interface must be able to reach the targeted devices.
 - The target device(s) must already be modelled within the Darktrace Threat Visualizer.
 - Responses from targeted device(s) must be observable by Darktrace.
 - Targeted device(s) must be able to accept and respond to the protocol used for identification.
- Darktrace / OT is provided with continuous, comprehensive visibility of OT-relevant network traffic. Mirrored traffic must be of suitable quality and free from duplication, fragmentary data, filtering, or unidirectional feeds that would inhibit effective analysis.
- Where non-IP or serial-based OT assets are observed via gateways or protocol converters, gateway configurations must remain stable so that device associations remain observable to Darktrace.
- OT-specific protocols (for example, Modbus, DNP3, S7, IEC-104, PROFINET, OPC, BACnet, and others supported by Darktrace / OT) must be visible through appropriate mirroring or ingestion points.
- Physical Darktrace appliances deployed in OT environments, including IMP units, must be installed in accordance with Darktrace guidance to ensure sufficient visibility across operational zones.
- Required visibility must be preserved following any changes to segmentation, routing, firewalling, or broader topology.
- Devices intended for OT-specific analysis must not be concealed behind encryption, tunneling, or other mechanisms that prevent visibility.
- Stable addressing and consistent communication patterns must be maintained to support long-term OT pattern-of-life behavioral modelling.
- Device, version, and protocol information required for Operational Alerts and OT Risk Management must be observable in traffic or supplied through integrations.
- Where risk datasets or external threat intelligence sources require outbound connectivity, suitable access must be provided or offline update mechanisms must be used.

- Darktrace / OT components must operate within expected capacity limits, with traffic volumes and composition kept within the performance specification of deployed Darktrace instances.
- The positioning of the Darktrace appliances, the selected interface and the configuration of the network and firewalls to allow a logical route for the connection to the target device(s), must be clarified to ensure queries can be successfully sent and received.
- Use of Autonomous Response in OT environments must include appropriate operational oversight due to the sensitivity of industrial equipment to connectivity disruption. Safeguards may include, where appropriate, human-confirmation mode operation, restricted device eligibility, or use of integrations. *Responsibility for ensuring controls remain appropriate lies with the client operator.*

7. Customer Responsibilities

All material outlined in section “5. Customer Responsibilities” of the Darktrace / NETWORK specification is applicable here.

Notwithstanding any additional responsibilities set out in any other section of this Product Specification, and without limitation, Customer’s responsibilities are to ensure that:

- Darktrace / OT is deployed in line with Darktrace best-practice recommendations and receives continuous, comprehensive visibility of OT-relevant traffic. Mirrored traffic is of suitable quality and free from duplication, filtering, or fragmentation.
- OT-specific protocols (for example, Modbus, DNP3, S7, IEC-104, PROFINET, OPC, BACnet) remain visible at capture points, and remain visible following any adjustments to segmentation, routing, VLANs, firewall policies, or broader topology.
- Physical appliances and virtual sensors deployed in OT environments are installed in accordance with Darktrace guidance, ensuring appropriate placement and sufficient visibility.
- For non-IP or serial-based OT assets observed through gateways or protocol converters, gateway configurations remain stable and preserve identifiers and protocol detail.
- Changes to OT network architecture (including segmentation, routing, or firewalling) are accompanied by timely updates to mirroring/ingestion so that coverage and analysis scope are preserved.
- The Active Device Identification module (if in use) is configured appropriately for the environment and used only with the consent of asset owners.
- Darktrace / OT integrations (such as asset-management platforms, vulnerability databases, lifecycle datasets, or compliance modules) are configured according to their documentation, with required access and connectivity maintained.
- Outbound connectivity required for third-party datasets is maintained, or offline update mechanisms are used.
- If Autonomous Response is enabled:
 - It is licensed and enabled globally (and, where applicable, in the appliance console).
 - Device eligibility and autonomy settings are configured appropriately (including tag-based eligibility and scheduling).
 - Reachability to all intended devices is maintained.

- Where RST packets cannot cross stateful boundaries, routing and/or firewall integrations are implemented.

Please refer to the Darktrace / NETWORK specification for further detail.

Additionally, Autonomous Response configurations account for the heightened sensitivity of OT devices to connectivity interruption. Actions that interrupt or restrict OT communication are not to be enabled unless the Customer has satisfied itself that such actions will not adversely impact operational processes.

- System health issues are monitored and resolved promptly.
- End-user access is managed securely and in accordance with best practice.

8. Considerations

All material outlined in “6. Considerations” in the Darktrace / NETWORK specification is applicable here.

Where source data is low quality, represents only a subset of OT network activity, is filtered, duplicated, asymmetric, delayed by external factors, or otherwise incomplete, the quality of Darktrace analysis will be impacted.

Physical Darktrace appliances deployed in OT environments, including ruggedized IMP units, must be positioned and configured in accordance with Darktrace guidance to ensure appropriate visibility while avoiding duplicate or fragmentary feeds. Changes to segmentation, routing, VLANs, firewall policies, or overall topology may alter capture paths and introduce visibility gaps until mirroring or ingestion is updated accordingly.

OT protocols must also be made visible at appropriate capture points to ensure Darktrace OT protocol analysis is possible. Where application-layer visibility is incomplete, the fidelity of detection and classification will be reduced.

OT capabilities must also be enabled on probes, vSensors, and other topology components intended to perform OT-specific Deep Packet Inspection. If these capabilities are disabled, only NETWORK-level analysis will occur. Darktrace / OT supports limited modelling of non-IP device types observed via gateways or similar technologies; devices not surfaced through compatible mechanisms cannot be enumerated or modelled.

Outputs that rely on historical analysis, such as IEC 62443-3-3 compliance reporting, may be incomplete immediately following deployment or after significant environmental changes.

External enrichment sources, including vulnerability and lifecycle datasets, may be limited or outdated when connectivity is restricted. Offline or alternative update mechanisms should be used where supported.

Use of Active Device Identification must be undertaken with the consent of asset owner(s), as some OT devices are highly sensitive to probing. While single-device mode carries lower risk than automated discovery modes, but a residual risk of disruption remains.

OT devices are more sensitive to loss of connectivity or unexpected interruption. Care must therefore be taken when enabling or configuring Autonomous Response in OT environments. Actions that restrict or terminate communication may have operational impacts. Where applicable, firewall/active integrations may be used to mitigate limitations of TCP-based actions (for example, against UDP-based traffic), but the responsibility to ensure such controls do not disrupt OT operations remains with the Customer.

9. Roles

All material outlined in “7. Roles” in the Darktrace / NETWORK Product Specification is applicable here. In addition, the following roles may be involved in OT specifically.

Customer Role	Responsibility
OT Security Lead/ Manager	Coordinates Customer teams focused on OT.
OT Security Specialist	Leverages OT Product functionality; monitors & actions output. Provides feedback on Product usage, input and outputs to other Customer & Darktrace roles

Darktrace / IDENTITY

- 1. Introduction..... 51
- 2. Product Features 52
 - 2.1. High Level Summary 52
 - 2.2. Darktrace Analysis 53
 - 2.3. Output 56
 - 2.4. Interface 63
 - 2.5. Reporting 64
- 3. Deployment Architecture 65
- 4. Implementation 65
 - 4.1. License Keys 66
 - 4.2. Deploying Modules 66
 - 4.3. External Alerting, Mobile App and Model Editing 67
- 5. Administration..... 67
 - 5.1. Administration of Darktrace / IDENTITY Modules 68
 - 5.2. De-Commissioning 68
- 6. Requirements 68
- 7. Customer Responsibilities 68
- 8. Considerations..... 69

1. Introduction

Darktrace / IDENTITY is a market leading cyber security solution that integrates with enterprise software and cloud platform solutions to bring visibility and threat analysis to critical systems. Extending Darktrace’s ActiveAI Platform beyond the physical enterprise network, each module brings the insight of the Cyber AI Analyst and Darktrace’s unique ‘pattern of life’ anomaly detection to enterprise software and cloud-based environments, monitoring user activity whether it originates inside the network or from remote locations.

This document is to be read in conjunction with the Darktrace Master Service Agreement which governs the usage of the Darktrace Product Offering. References to “Customer” throughout this document should be read to refer to the entity that is the owner of the software subscription and is ultimately responsible for its operation, whether as end-user or service provider.

Darktrace / IDENTITY Subscriptions are sold according to the Usage Metrics bands set out in the relevant Product Order Form.

Darktrace / IDENTITY was previously named Darktrace DETECT/Apps, Darktrace RESPOND/Apps, Darktrace DETECT/Zero Trust, and Darktrace RESPOND/Zero Trust (excluding Zero Trust/Netskope and Zero Trust/Zscaler ZIA/ZPA). Darktrace / IDENTITY that is implemented into cloud-hosted environments may be referred to in exceptional circumstances as Darktrace / CLOUD (Identities). This specification applies for Customers who have purchased the Product under any of the above naming conventions.

2. Product Features

2.1. High Level Summary

Darktrace / IDENTITY modules retrieve activity data from third-party SaaS and enterprise software platforms.

Activity data is processed, parsed, and analyzed to identify users or other entities which are active within the third-party platform. Each user or entity identified is then evaluated for behavior outside of a “normal” state, created by an ongoing, real-time analysis of audited activity and other supplied data events across the enterprise network environment. The system profiles individual users or entities based upon their activity, and the activity of those it deems “peers” due to similar behavioral activity (“modelling”).

Alerts are raised when specific criteria, or a minimum threshold of unusual behavior is met. The output is surfaced to operators in the Darktrace / IDENTITY Console and Darktrace Threat Visualizer interface for investigation and resolution. Some Darktrace / IDENTITY modules can also respond to unusual activity (“Autonomous Response”) by modifying the users, entities, or relevant configuration settings within the third-party platform.

2.1.1.Data Retrieval and Ingestion

Darktrace / IDENTITY modules retrieve log-based activity data from third-party platforms; this is typically achieved via APIs or export formats provided by the third-party platform. Due to the variation in APIs, authentication methods, and SDKs, provided by each separate platform, the methodology for ingestion will differ between Darktrace / IDENTITY modules. The method of authentication and retrieval is outlined in the relevant documentation provided for each module.

Delays may be incurred where the external platform does not make events available to the Darktrace module for processing and analysis within the expected time frame. Delays of this nature are the responsibility of the associated third-party platform vendor. Latency between event occurrence and when it was made available to the module are indicated in the event metadata within the Threat Visualizer.

A subset of modules will also retrieve contextual enrichment data on users, files, and other relevant entities. The data available and the method by which it is retrieved are outlined in the relevant documentation provided for each module.

Darktrace offers a suite of Threat Intelligence and Telemetry integrations where data may be retrieved from other compatible security tools or OSINT sources. The retrieval method for each integration is detailed in the corresponding integration documentation. These alternative inputs (“Threat Intelligence Integrations” & “Telemetry Integrations”) are configured on the Darktrace Threat Visualizer “System Config” page.

Darktrace also provides a REST API for automation and a subset of relevant data input.

2.1.1.1. “Open” Ingestion

Data ingested by Darktrace / IDENTITY modules is retrieved directly from the audited activity logs of the relevant third-party platforms - returned information is therefore limited to the events that each vendor chooses to audit, and the data recorded as part of those audit log entries. Variation in both the events available to each Darktrace / IDENTITY module, and the level of detail, are expected.

Darktrace / IDENTITY modules operate in a passive, “open” ingestion mode - modules do not actively request specific events or place restrictions on the data retrieved from the third-party platform. The range of events available to the module will therefore vary both between each monitored platform and between each client.

Some third-party platforms may also restrict the events and APIs available behind license requirements. Darktrace will outline where possible the minimum license requirements required for basic real-time monitoring operation.

Exceptions exist where the resource types to be monitored are pre-defined during configuration (e.g. Salesforce, Google Workspace), or where individual event types must be removed due to excessive volume.

2.2. Darktrace Analysis

Darktrace analysis is multi-stage. The following outline highlights key elements of the process but is not intended to be an exhaustive list of all analysis performed during operation of the Darktrace platform, Threat Visualizer interface, or any underlying components.

Where source data is low quality, represents only a subset of all activity, is delayed by external factors, or otherwise is incomplete, the quality of Darktrace analysis may be severely impacted. Analysis may also be impacted by overloading, such as that observed when event throughput significantly increases beyond scoped levels.

2.2.1. Darktrace / IDENTITY Event Inspection

Darktrace / IDENTITY modules will retrieve activity log information made available by third-party APIs (or other comparable methods) for analysis. Event and activity data is then parsed into a series of standard metadata fields and analyzed to identify the type of activity, any entities involved (e.g., users, files, or other resources), and to extract key information.

The output of this analysis is passed to the Darktrace model engine, to the core Darktrace “classifier” engine, and made available for display in the user interface and any other platform components such as the Cyber AI Analyst.

2.2.2. Darktrace ‘Pattern of Life’ Analysis and Classification

Network events, user activity events, connection data and any other configured inputs are subject to Darktrace ‘pattern of life’ analysis. Darktrace will create individualized behavioral profiles for the network entities it observes and surface activity which is considered inconsistent with the expected norm. This baseline is derived from - but not limited to - an analysis of the behavior of the individual entity, analysis of one or more clusters created from similarly behaving entities, and many variable factors such as time or communication protocol. The ‘pattern of life’ data is continually updated in real time, and reflects the data that has been received, with a greater weighting to more recent data.

Analysis is performed in the “classifier” stage by a multitude of classifiers. This core analysis applies many approaches including unsupervised Machine Learning techniques such as Bayesian Meta-

classification, techniques derived from graph theory and network analysis such as Node/Graph Centrality, approaches derived from statistical analysis such as Spectral Clustering and many other techniques. The previous examples of utilized techniques provided are for illustrative purposes and should not be considered exhaustive.

Darktrace does not offer any capability to access underlying behavioral models or classification output. Darktrace provides the ability to utilize the output of this analysis using the *models* framework. Darktrace will create alerts to indicate anomalous activity which will be inserted into the “event log” within the Darktrace Threat Visualizer interface of the corresponding device, user, or entity. These unusual activity “notices” are also consumed by Darktrace Real-Time Detection models.

2.2.3. Darktrace Real-Time Detection Models Framework

The Darktrace Real-Time Detection model engine is the logical framework within which events and the output of ‘pattern of life’ analysis is evaluated. Both analysis components described above submit data to the model engine for evaluation.

A model is used to define a set of conditions which, when met, will alert the system to the occurrence of a particular event or chain of anomalous behavior. The models framework leverages both the underlying ‘pattern of life’ detection and outputs from Darktrace Deep Packet Inspection, telemetry inputs, Darktrace / IDENTITY and Darktrace / CLOUD modules. Output from the complex anomaly framework is available in accessible, building block format and can be combined with simple conditions and logical expressions to create tailored activity detection.

Custom models can be defined by operators of the Darktrace Threat Visualizer to meet any criteria; these are created and maintained in the Threat Visualizer “Model Editor” interface. Care should be taken when defining custom models to ensure that the number of alerts produced is not excessive and does not impact the system’s ability to function, or otherwise degrade service. Similarly, modification of existing default models should not result in excessive activity or alter the logic to the extent that alerts criteria can no longer be met. The responsibility to ensure that models created or edited remain within these reasonable expected boundaries lies with the Customer.

Default Darktrace models are focused on ‘pattern of life’ anomaly detection, potentially malicious behavior and optional compliance issues. Darktrace will periodically update these standard supplied models - customers with Call-Home or Darktrace-hosted cloud instances will receive updates automatically, clients without automatic updates will receive all applicable model updates when Darktrace Threat Visualizer software is updated. The conditions for auto-update eligibility are described in the relevant product documentation.

Models may also be used to trigger actions within the Darktrace Threat Visualizer platform; the output of the Darktrace Real-Time Detection model engine is described below.

2.2.3.1. Autonomous Response Models

In default operation, Darktrace / IDENTITY Autonomous Response responses are triggered by model alerts from a specific subset of Darktrace models, categorized as Autonomous Response (“Antigena”) models.

Darktrace Autonomous Response models may directly look for specific behavior, or for indicators identified by other models operating within the real-time Detection environment. Darktrace / IDENTITY models fall typically into the second category - “meta-models” - which are triggered by an alert of another Darktrace Real-Time Detection model.

Any modification to an underlying Real-Time Detection model which results in increased or reduced model alerts will subsequently impact Darktrace / IDENTITY Autonomous Response models and actions.

Any under-activity or over-activity of the corresponding Darktrace / IDENTITY Autonomous Response model as a result of this type of modification is the responsibility of the Customer.

Darktrace also provides the ability to limit actions created by these models to a minimum score threshold.

2.2.4. Cyber AI Analyst analysis

Darktrace Cyber AI Analyst performs a meta-analysis upon the previous layers of analysis described. Please refer to Cyber AI Analyst in the Product Agnostic section below.

2.2.5. Cross-Capability / Cross-Coverage Area Analysis

Where other Darktrace ActiveAI Security Platform coverage areas are deployed, cross-platform analysis is performed by multiple components. Links are created for analysis purposes between network entities and entities modeled by other Darktrace Real-Time Detection components such as, for example, user entities created by Darktrace / IDENTITY module monitoring known to be associated with a given network entity, or credential entries observed by Darktrace / NETWORK monitoring also observed as part of Darktrace / ENDPOINT.

Darktrace Cyber AI Analyst will also link together entities across different coverage areas in the creation of Cyber AI Analyst Incidents and Cyber AI Analyst Incidents Events. Darktrace Cyber AI Analyst may also retrieve additional contextual data during investigation from other components if deployed, such as the retrieval of associated emails from Darktrace/Email. Please refer to “Cyber AI Analyst” below.

The above provides illustrative examples of collaborative scenarios between platform components for reference but is non-exhaustive.

2.2.6. “User” Device Entities

As part of the analysis process, Darktrace will model individualized behavioral profiles in the form of “user” device entities.

Each “user” device represents a distinct, unique actor that has performed a *meaningful* action in the third-party platform. This is defined as any successful login, or any “outbound” event, where outbound indicates the user or entity was the actor who performed the action. This is distinct from “inbound” actions, where the user or entity was observed within the context of another user modifying them in some way but was never seen performing any actions themselves.

Examples of outbound actions may include a user accessing files, creating API keys, or modifying a resource in some way. In this definition, a “resource” is any modifiable element, including but not limited to a group, a file, a virtualized computing resource, a password policy, or another user.

Failed login events alone will not create user entities but will be added to user entities if created already by a meaningful action.

Where a user or entity is only observed performing an event which does not meet the criteria for the creation of a “user” device - and no prior user device exists for them - the event will be aggregated into a generic log for the Darktrace / IDENTITY module. Behavioral analysis and Darktrace Real-Time Detection will be severely limited for this aggregate entity.

Many third-party platforms permit external entities to influence resources. Similarly, many third-party platforms and systems operate “service” accounts or system-owned entities that may alter internal resources. Therefore, Darktrace creates no expectation that the count of user device entities created from activity monitoring will directly align with license or seat counts negotiated with the third-party platform.

Darktrace provides mechanisms to limit the scope of entities modeled in this way.

2.2.6.1. Restricting “User” Device Entities

Organizations may wish to restrict visibility over the event data retrieved and modelled by Darktrace / IDENTITY modules due to regional monitoring restrictions, multi-tenant environments, or other compliance policies. Darktrace offers both user-level and ingestion-level restrictions to limit the processing of this data. Typically, the filtering of data is performed *after it is retrieved* from the third-party platform due to the open ingestion model outlined above.

The “Activity Filter” and “Group Activity Filter” (*limited modules only*) settings restrict the creation and modelling of users and groups respectively to those that meet or do not meet a specified criterion. Users which take actions that impact explicitly monitored users, or are directly actioned by monitored users, will still be modeled if they do not match the criteria to ensure that the two-way interaction is captured. This is the recommended operating mode to ensure full visibility of relevant information for operators.

If it is necessary to ensure that information about these unmonitored users is not preserved, Darktrace provides mechanisms to collapse all users outside the scope of monitoring to a single entity. Behavioral analysis and Darktrace Real-Time Detection will be severely limited for this aggregate entity.

2.3. Output

Darktrace / IDENTITY modules can possess both Darktrace Real-Time Detection and, where applicable and licensed, Darktrace Autonomous Response capabilities. The output of these two capabilities are outlined individually.

2.3.1. Real-Time Detection

Darktrace / IDENTITY Real-Time Detection outputs activity-based metadata for users and entities, formatted into consistent categories and with the most relevant information pre-extracted. As noted above (refer to 1.2.1.1 “Open” Ingestion), each module will retrieve all possible relevant event activity from the third-party platform - the events retrieved and surfaced will therefore differ between monitored platforms and between client deployments.

Darktrace / IDENTITY Real-Time Detection will typically surface activity data, results of ongoing ‘pattern of life’ analysis (including notification of unusual activity), model events, Darktrace Cyber AI Analyst alerts, contextual data retrieved from configured integrations (optional), and any other log-based event data created or processed by the Darktrace system. The platform may also output alerts about overall system health and component health.

Deploying a Darktrace / IDENTITY module will provide access to the Darktrace / IDENTITY Console, a specialized interface for investigating SaaS and Cloud activity. The Console provides access to global maps of activity, detailed logs of user activity and visualization of anomalous chains of behavior. Where Darktrace / NETWORK or another comparable Darktrace ActiveAI Security Platform coverage area is also present, activity processed is also surfaced in the main Darktrace Threat Visualizer interface. Data is optionally available through the Darktrace Mobile App, compatible output formats, or the Darktrace Threat Visualizer REST API.

2.3.1.1. Processed Event Metadata

The output of Darktrace / IDENTITY Real-Time Detection is displayed in the Darktrace Threat Visualizer “Advanced Search” interface. Data retention for this output is on a rolling basis and is dependent upon a

number of factors such as hardware capability, event volume and other operational components which contribute entries to Advanced Search. Most deployments can expect around 30 days retention.

2.3.1.2. Alerts

Darktrace / IDENTITY Real-Time Detection produces two primary alert types: Darktrace Model Alert (*formerly known as “Model Breach” alerts*) and Darktrace Cyber AI Analyst Incidents. Darktrace Model Alerts and Darktrace Cyber AI Analyst Incidents can be investigated in the Darktrace / IDENTITY Console, the main Threat Visualizer interface, investigated in the Darktrace Mobile App, or exported to a compatible alert consumer. Alerts are categorized for priority and filterable by multiple factors, allowing for customization of alerts displayed in Darktrace interfaces and those exported to external tools.

A secondary alert type - Darktrace System Status alerts - may also be generated where the Darktrace platform is experiencing degraded service due to health issues, invalid input, or where it is necessary to highlight changes to system administrators. Alerts include details of the originating host, the severity of the event, and relevant links to investigate or resolve the issue. Notifications are sent when a system event becomes active and can optionally be sent on resolution.

Darktrace supports alert export in both industry-standard generic forms such as Syslog or Email and custom integrations with common tools such as Splunk and ServiceNow. Alert outputs (“Workflow Integrations”) are configured on the Darktrace Threat Visualizer “System Config” page. The details included in each external output may vary due to third-party restrictions on content length or supported formats.

2.3.1.2.1. Model Alerts

Darktrace model alerts are created as a result of Darktrace / IDENTITY Real-Time Detection models; when conditions for a model are met, a model alert can be created in addition to other possible model actions. Darktrace model alerts contain details of the conditions that were satisfied in order to trigger the alert, the entity which met those conditions (for example, a device, or another model alert) and a description with recommended action points.

Model alerts will only be created for users who are included in monitoring.

2.3.1.2.2. Darktrace Cyber AI Analyst Incident Alerts

Darktrace Cyber AI Analyst Incident alerts are created when Darktrace Cyber AI Analyst identifies activity considered significant enough to highlight to operators. Darktrace Cyber AI Analyst Incidents may contain a single “Cyber AI Analyst Incident Event”, or multiple linked findings aggregated together. External alerts are created when a new “Cyber AI Analyst Incident Event” is created, which may be associated with an existing “Cyber AI Analyst Incident”, or form an independent, new “Cyber AI Analyst Incident”. In the former case, mechanisms to identify the relationship between Cyber AI Analyst Incident Events are provided in the output.

In the Darktrace / IDENTITY Console, Darktrace Threat Visualizer and Darktrace Mobile App, the output of Darktrace Cyber AI Analyst analysis is aggregated on a per-Cyber AI Analyst Incident basis. New Cyber AI Analyst Incident Events added to existing Incidents will not produce independent alerts in these interfaces and are instead displayed as part of the existing Incident. Darktrace Cyber AI Analyst Incidents in these interfaces contains details of the activity identified, the Darktrace model alerts that triggered the initial Cyber AI Analyst investigation, the entities which performed the unusual activity, (for example, devices), the investigation steps Cyber AI Analyst performed, why activity was linked together (if multiple Cyber AI Analyst Incident Events), and a human-readable summary of the finding.

2.3.2. Autonomous Response

Darktrace / IDENTITY Autonomous Response actions can be taken when Darktrace / IDENTITY Real-Time Detection identifies significantly anomalous user behavior, when Darktrace identifies activity which contravenes a compliance policy, when manually triggered by an operator, or as a result of any other custom criteria defined in a Darktrace model. Darktrace / IDENTITY Autonomous Response is only available for a specific subset of supported Darktrace / IDENTITY modules.

If the current system configuration permits autonomous actions - or the action was triggered manually - the action will be created in *active* state. This created action will appear on the Darktrace / IDENTITY Console Response Actions Page under “Active” actions, in the Darktrace Threat Visualizer Response Actions Page “Platform” tab, and in the Darktrace Mobile App. If alerts for Autonomous Response actions (“*created*” state) are configured, an alert will also be sent to configured outputs.

If Darktrace / IDENTITY Autonomous Response is compelled to request human confirmation due to the current system configuration, the action will be created in *pending* mode. This created action will appear on the Darktrace / IDENTITY Console Response Actions Page (and locations outlined above) under “Pending” action until a user confirms that Darktrace / IDENTITY can take action. A notification requesting that a human user approve the action is sent to the Darktrace Mobile App and any configured Darktrace Autonomous Response alert outputs (“*pending*” state). Once confirmed by an operator, the action will move to “Active” as before.

Read-only access to view Darktrace / IDENTITY Autonomous Response actions is not restricted for users with appropriate permissions to view Darktrace / IDENTITY Real-Time Detection data.

Darktrace / IDENTITY operates an “opt-out” approach to Autonomous Response eligibility. Please refer to “1.2.8 Autonomous Response Eligibility” below for further information.

2.3.2.1. Pending vs Active Actions

Darktrace / IDENTITY Autonomous Response actions created as a result of Darktrace Real-Time Detections are initialized in two possible states: *active* and *pending*.

An *active* Darktrace / IDENTITY Autonomous Response action permits Darktrace / IDENTITY to take the responses defined by the action configuration. For example, to log a user out or block a specific IP via the third-party platform API or supported integration method. As manual Darktrace / IDENTITY response actions are created directly by an operator, actions are always created in the *active* state.

Pending actions are those which Darktrace / IDENTITY Autonomous Response wishes to take but must require approval before it can act automatically; pending actions may therefore be referred to as “human confirmation mode” actions. Pending actions do not cause modifications in the third-party platform until approval is granted in some format, at which the point the action becomes “active”.

Darktrace provides mechanisms to alert users on the creation of pending actions requiring human approval to proceed. These notifications are surfaced in the main Threat Visualizer interface (*excludes Darktrace / IDENTITY Console*), in the Darktrace Mobile App, and in compatible alert outputs.

Please refer to “1.2.9 Darktrace / IDENTITY Response Autonomy” for more information about these system conditions

The alternative terminology “Human Confirmation” or “Human Confirmation Mode” may also be used, respectively, when referring to pending actions or a system state where pending actions are created.

2.3.2.2. Action States

Darktrace Autonomous Response actions proceed through several states during their lifecycle. For example, an action is automatically created that *requires confirmation* (“pending”), it is then *confirmed* by a user. Another user *extended* the action, then *cleared* it. Finally, it was *reactivated* by another user, before *expiring*.

There are four key states (“pending”, “active”, “cleared”, “expired”) and two additional states (“extended”, “reactivated”):

- The first two states - pending and active are outlined above
- Cleared actions are those which have been manually ended by an operator. Clearing informs Darktrace to cease the action, and to suppress the combination of Darktrace / IDENTITY Autonomous Response action and model alert conditions for the remainder of the action's set duration.
- Expired actions are those which are historic, regardless of their state (pending, active) before the action period passed.
- Reactivated actions are those which were made active again by an operator.
- Extended actions are those which are currently active and have had their duration manually extended by an operator.

Actions can be extended, reactivated or cleared from the Response Actions page and the Darktrace Mobile App.

2.3.2.3. Actions in Default Operating Mode

Default operating mode (*default operation*) is here defined as Darktrace / IDENTITY Autonomous Response operating in a state where autonomous (*active*) actions can be taken automatically against eligible users on supported platforms, when triggered by a default Darktrace / IDENTITY Autonomous Response model. Autonomous actions are actions active without human intervention.

Default operation presumes that Darktrace / IDENTITY Autonomous Response is enabled globally on the Darktrace ActiveAI Security Platform instance, licensed for the appropriate module, and the module has been granted the appropriate permissions in the third-party platform to enact actions. In some cases, it may be necessary to re-authenticate the module to grant the necessary permissions. Please refer to the corresponding documentation for each module for further guidance.

The conditions for the creation of action in any state in default operation are:

1. A Darktrace / IDENTITY Autonomous Response model which has its logical criteria sufficiently met.
2. The prior criteria met by a user or IP Address which has not been explicitly excluded from actions.
3. A Darktrace / IDENTITY Autonomous Response action (inhibitor) configured within the triggered model to define the action taken, where:
 - The prior criteria are not met by an entity which is “immune” to the inhibitor in the module configuration.
4. A system operating mode (schedule) which does not prevent the creation of a Darktrace Autonomous Response action at the given time.

The conditions for an *active* action in default operation are:

5. A Darktrace / IDENTITY Autonomous Response model which has its logical criteria sufficiently met.
6. The prior criteria met by a user or IP Address which has not been explicitly excluded from actions.
7. A Darktrace / IDENTITY Autonomous Response action (inhibitor) configured within the triggered model to define the action taken, where:
 - The prior criteria are not met by an entity which is “immune” to the inhibitor in the module configuration.
8. The model in question is configured to *permit* or *force* autonomous actions.
9. The system operating mode permits the creation of an active Autonomous Response action at the given time for the given device.

Manual actions triggered by an operator are not subject to the conditions outlined above; these actions are always active on creation.

2.3.3. Darktrace / IDENTITY Actions

Darktrace / IDENTITY response actions can be triggered by Darktrace / IDENTITY Autonomous Response models, by custom models possessing the required configuration, or manually by an operator.

The actions Darktrace / IDENTITY takes in response to a trigger will vary based upon the action configuration and the third-party platform the trigger user has been observed upon. The types of response actions taken are henceforth referred to as “inhibitors”, as they *inhibit* a specific form of behavior or activity.

The range of inhibitors available differ between Darktrace / IDENTITY modules.

2.3.3.1. Manual Darktrace / IDENTITY Response Actions

Darktrace provides mechanisms to create Darktrace / IDENTITY response actions manually. These actions can be created by a user with appropriate permissions from the Darktrace / IDENTITY Console, the main Threat Visualizer interface, or from the Darktrace Mobile App.

Darktrace provides the option to enforce a free-text justification which must be submitted by the end-user when the action is created.

Manual actions are automatically created in an “active” state and do not require human activation. The type of action taken (*inhibitor*) is selected during creation. The responsibility to select an appropriate inhibitor lies with the Customer.

Where a user or targeted IP Address is excluded from the specific action inhibitor, or from all actions, triggered manually by an operator, this will be indicated in the action status. In this case, an active action will be created but the inhibitor will be prevented from taking effect.

2.3.3.2. Model-Triggered Darktrace / IDENTITY Autonomous Response Actions

In default operating mode, Darktrace / IDENTITY Autonomous Response actions are triggered by models. This trigger is referred to as a “model action” - a system action taken in response to a specific model criterion being met.

The creation of a Darktrace Autonomous Response action is a model action which exists by default on applicable Darktrace-created models created for the purposes of Autonomous Response. When the criteria for these models are met, Darktrace will invoke all model actions applied to the model, triggering a Darktrace / IDENTITY Autonomous Response action as a result. In this context, “trigger all model actions”

refers to all possible model actions applied to the model (for example, “tag device” or “create alert”). It does not refer to triggering all *inhibitors* applied to the model.

Operators may also add the Darktrace Autonomous Response model action to modified or newly created custom models. The responsibility for selecting an appropriate action inhibitor (see below), and for any actions created as a result of a custom model configuration, lies with the Customer.

2.3.3.3. Darktrace / IDENTITY Autonomous Response Inhibitors

The type of action Darktrace / IDENTITY Autonomous Response takes is referred to as an inhibitor. The suite of available inhibitors differs between each platform monitored by Darktrace / IDENTITY - models therefore accept multiple Darktrace / IDENTITY inhibitors to ensure an appropriate action exists across all relevant platforms. The supported inhibitors for each platform can be found in the relevant Darktrace documentation.

In default operation, when the criteria for a Darktrace / IDENTITY Autonomous Response model are met, Darktrace / IDENTITY will attempt to create an action corresponding to the selected inhibitor for the combination of trigger user and platform. The inhibitor is defined in the model directly (see “1.2.3.4.2 Model-Triggered Darktrace / IDENTITY Autonomous Response Actions”).

Manual actions require the inhibitor to be selected during creation. This is applicable to actions created through the Darktrace / IDENTITY Console, main Threat Visualizer, and the Darktrace Mobile App.

For actions created as a result of a user-created or user-modified model, Darktrace / IDENTITY Autonomous Response will attempt to create an action corresponding to the inhibitor as defined in the model (see “1.2.3.4.2 Model-Triggered Darktrace / IDENTITY Autonomous Response Actions”) for the trigger user and platform.

Darktrace / IDENTITY allows individual inhibitors to be disabled (*specific modules only*). Users and IP addresses can also be made “immune” on a per-inhibitor basis.

2.3.3.3.1. Modification and Restriction of Inhibitors

Typically, Darktrace / IDENTITY Autonomous Response will take the action (“inhibitor”) selected in the model action in response to the Darktrace / IDENTITY Autonomous Response model criteria being met. However, there are exceptions to this behavior:

- Where the user who met the model criteria was observed on a platform which does not support Autonomous Response. In this case, no actions are created.
- Where Darktrace / IDENTITY possesses some of the necessary permissions to take actions but is unable to target all user types, the action may be downgraded to a supported capability (*relevant modules only*).

For example, if Darktrace / IDENTITY does not have the appropriate user role to target Microsoft 365 “Global Admin” users with a “disable user” inhibitor, it will attempt to downgrade to a “force logout” action. This will be noted in the action status on the Response Actions page.

Downgrading can be disabled in the relevant module configuration.

- Where the targeted user is configured to be “immune” from the specific inhibitor, or globally from actions. In this case, an active action will be created but the inhibitor will be prevented from taking effect. This will be noted in the action status on the Response Actions page.

- Where the targeted IP Address is configured to be “immune” from the specific inhibitor, or globally from actions. In this case, an active action will be created but the inhibitor will be prevented from taking effect. This will be noted in the action status on the Response Actions page.
- Where the targeted entity is of an unsupported type (as outlined in the relevant documentation). Darktrace / IDENTITY may attempt to take action - an active action will be created - but be prevented from succeeding. This will be noted in the action status on the Response Actions page.
- Where the triggered inhibitor is “Block IP”, and the model name meets the criteria of the “Block IP without Force Logout Models” field in the module configuration. In this case, “Block IP” will be applied but only in a partial format.

2.3.4. Autonomous Response Actions Initiated by Cyber AI Analyst

Cyber AI Analyst can now initiate action creation - this functionality is optional and disabled by default.

Actions are created as a result of Cyber AI Analyst Incidents. Cyber AI Analyst links related devices and unusual activity together to create an overall incident structure. If Cyber AI Analyst detects that a link between two entities does not have a corresponding Autonomous Response action - and an inhibitor is available which would suitably target the activity exists - it can now create an action to target this behavior.

2.3.5. Autonomous Response Eligibility

Darktrace / IDENTITY operates an “opt-out” approach to Autonomous Response eligibility. All user entities detected by Darktrace / IDENTITY modules with Autonomous Response capabilities are considered eligible for Autonomous Response actions.

Eligibility is highly configurable - both users and IP addresses (relevant to the “Block IP” inhibitor) can be made *immune* from actions on a global-level, module-level or inhibitor-level.

Where the Darktrace / IDENTITY module is active on a Unified View instance (*specific modules only, advanced operating mode*), immunity must be configured on the corresponding subordinate master where the user entity is modelled.

2.3.6. Darktrace / IDENTITY Response Autonomy

Darktrace / IDENTITY Autonomous Response can create actions in an “active” or “pending” state, were pending actions require human approval before any direct action is actually taken. The state actions are created in - referred to here as “autonomy” - is defined by the following factors and conditions.

2.3.6.1. Manually Triggered Darktrace / IDENTITY Response Actions

Manual actions are always automatically created in an “active” state.

2.3.6.2. Model-Triggered Darktrace / IDENTITY Autonomous Response Actions

The status of actions created as a result of a Darktrace model alerts (both default Darktrace / IDENTITY Autonomous Response models, and those user-created or user-modified) is defined by a combination of granular eligibility controls.

2.3.6.2.1. *Model-Based Autonomy*

Darktrace / IDENTITY Autonomous Response autonomy can be configured on a per-model basis to meet the varied requirements of each organization. This autonomy is controlled in the model configuration: models may “force human confirmation”, “force autonomous action”, or “permit autonomous action”. Actions created by models with a “force” state will be placed in a *pending* or *active* state respectively on creation, regardless of any other configuration. An action created as a result of the model with “permit autonomous action” will attempt to create an *active* action if permitted by other configuration factors but will defer to *pending* if overruled.

2.3.6.2.2. *Time-Based Autonomy*

Darktrace also provides a seven-day, hourly timetable (“Response Schedule”) which allows blocks of autonomous action or enforced human oversight to be scheduled. The schedule is applied across Darktrace / NETWORK, Darktrace / ENDPOINT and Darktrace / IDENTITY.

The schedule is applied to actions created by models with “permit autonomous action”. Actions created manually, or by models in a “force” state (“force human confirmation”, “force autonomous action”), are not impacted by the schedule state.

Localization is not available for Darktrace / IDENTITY module responses. If “local subnet time” is active, the schedule will apply Darktrace / IDENTITY actions in UTC.

2.3.6.2.3. *Cyber AI Analyst-Triggered Actions*

Actions created as a result of a Darktrace incident can “force human confirmation”, “force autonomous action”, or “permit autonomous action”. This setting is set on the Darktrace System Config page and applies to all actions created by Darktrace Cyber AI Analyst.

Actions created when this setting is in a “force” state will be placed in a *pending* or *active* state respectively on creation, regardless of any other configuration. An action when set “permit autonomous action” will attempt to create an *active* action if permitted by other configuration factors but will defer to *pending* if overruled.

2.4. Interface

The primary user interface of the Darktrace platform is the Threat Visualizer. The Threat Visualizer interface provides access to all major Darktrace user interfaces, consoles and product views - it contains both investigation interfaces and administration interfaces. A list of the user interfaces that comprise the Threat Visualizer - and detailed information on how to operate these interfaces - is provided in the relevant technical documentation.

The Darktrace / IDENTITY Console is a purpose-built investigation interface for Darktrace / IDENTITY that forms part of the Threat Visualizer.

- If Darktrace / IDENTITY is the only Darktrace coverage area, only the Darktrace / IDENTITY Console and a subset of relevant Threat Visualizer pages will be accessible.
- Where Darktrace / NETWORK or another comparable Darktrace ActiveAI Security Platform coverage area is also present, the Darktrace / IDENTITY Console will be available alongside the main Threat Visualizer interface. Darktrace / IDENTITY data will be surfaced in both views.

Both the Darktrace / IDENTITY Console and main Threat Visualizer interface include user customizable and filterable alert trays, a visualization of ongoing device activity over time, and access to the output of Darktrace analysis in both log and summary format. Alternative investigation views, access to review detailed metadata, report generation, and the ability to edit and create custom models (as described above) are also provided within these interfaces. Operators can also review the detected users observed by Darktrace / IDENTITY Real-Time Detection (including any relevant contextual data), configure system settings, deploy integrations, review system health information, perform user management and other administrative tasks.

Darktrace / IDENTITY Autonomous Response actions are recorded in the Response Actions window within the Darktrace / IDENTITY Console. This view restricts all actions to only those taken by Darktrace / IDENTITY modules. A per-user filtered view is also offered from each user profile.

Where Darktrace / NETWORK or another comparable Darktrace ActiveAI Security Platform coverage area is also present, response actions for Darktrace / IDENTITY are also surfaced in the main Darktrace Threat Visualizer Response Actions window. Here, Darktrace / IDENTITY actions are located under “Platform”.

Users with appropriate permissions may alter the state of an action from this view (clear, activate, extend, or reactivate). Darktrace / IDENTITY Autonomous Response actions are also surfaced in, and can be modified from, the Darktrace Mobile App.

Darktrace offers alerting when Darktrace / IDENTITY response actions are created or change state; notifications are surfaced in the Threat Visualizer interface, in the Darktrace Mobile App, and sent via compatible alert outputs.

Darktrace / IDENTITY modules display a current “state” indicating whether actions could be successfully applied, whether the user was immune in some format, or whether the action has subsequently expired.

2.5. Reporting

Darktrace offers both manual and automatic PDF report generation, scheduled via the Darktrace Threat Visualizer “System Config” page. Details of reporting formats offered is provided on the Darktrace Customer Portal and relevant documentation. Darktrace reserves the right to alter the content of reports offered to align with changing Product and Service Offering.

The Threat Visualizer “Audit Log” records changes made by operators such as model alert acknowledgement; the audit log can be exported via Syslog for extended retention. User interactions with Darktrace / IDENTITY Autonomous Response actions which alter the state (e.g. “activate”, “extend”, “clear”) are also recorded in the action history, accessible from the Darktrace Response Actions window. Users can also be compelled to provide a free-text justification when an action state is altered, which is also displayed where applicable.

The output of Darktrace / IDENTITY data retrieval and analysis is displayed in the Darktrace Threat Visualizer “Advanced Search” interface. Combined output from this processing, from Darktrace analysis, and from any actions performed by the platform automatically (such as tagging as a result of a model) are combined into logs which are displayed for each user.

Data retention for metadata output and combined log data is on a rolling basis and is dependent upon a number of factors such as hardware capability, event volume, and other operational components which contribute data to the platform. Most deployments can expect around 30 days retention of Darktrace / IDENTITY metadata and general device/user activity data. Automatic removal of older event log data does not affect the storage or training of the machine learning ‘pattern of life’ data.

User entities observed by Darktrace / IDENTITY Real-Time Detection to be recently active in monitored third-party platforms are detailed on the Darktrace / IDENTITY Console “Profiles” page. Where Darktrace

/ NETWORK or another comparable Darktrace ActiveAI Security Platform coverage area is also present, user entities will also be reported on the “Device Admin” page. Metrics regarding data and event throughput are rendered on the Threat Visualizer “System Status” page.

Darktrace also provides a REST API for automated retrieval of a subset of data output.

3. Deployment Architecture

Darktrace / IDENTITY modules can operate on Darktrace physical or cloud-hosted master instances, or on Unified View instances in select cases. Internet connectivity to the relevant endpoints - both those associated with the third-party platform, and with relevant Darktrace services - is necessary for data retrieval. Darktrace / IDENTITY modules only operate on physical probe instances, vSensors, osSensors, or other topology components referenced in the Darktrace / NETWORK specification.

The relevant combination of Darktrace / IDENTITY modules will vary depending on the customer technical stack and what they currently have in their environment. The list of currently available Darktrace / IDENTITY modules are as follows:

- Asana
- Box
- CloudFare
- DropBox
- Duo
- Egnyte
- Google Workspace
- HubSpot
- JumpCloud
- Microsoft 365
- Okta
- Salesforce
- Slack
- Zoom

The applicable Darktrace / IDENTITY modules may vary during the lifetime of a Darktrace deployment should new tools be introduced to the Customer’s architecture. It is the responsibility of the client to ensure that Darktrace maintains visibility over all traffic in the event of additions that will fall outside the scope of the initial deployment design.

Each deployment requires a minimum of one “*Master*” instance to provide the capabilities outlined above. Deployments should be structured so that a single instance - Master or Unified View - sits at the top of the deployment topology to operate the Threat Visualizer interface and other relevant components.

For further details on the relevant Topology Components and Topology Roles, please refer to the / NETWORK Product Specification.

4. Implementation

The Customer Portal contains technical guides for implementing the various Darktrace / IDENTITY modules. Customers should follow the best practices set out in these guides when implementing applicable modules.

4.1. License Keys

Darktrace / IDENTITY will only be available when a valid License Key is configured on the Threat Visualizer's System Config page. In Unified View environments, the License Key should propagate from the UV master instance to the subordinate Master instances unless specifically configured otherwise. Once the license is entered, Customer will get access to all available Darktrace / IDENTITY modules. Adding the License Key will also activate Autonomous Response for the modules that have this feature available. Instructions on adding a License Key can be found in the relevant documents in the Customer Portal.

4.2. Deploying Modules

Deploying one or more Darktrace/ IDENTITY modules will provide access to the SaaS Console, a specialized interface for investigating SaaS and Cloud activity. The console is powered by the Cyber AI Analyst and Darktrace's unique 'pattern of life' anomaly detection; each element is purpose built for monitoring and analysis in these environments whilst maintaining existing workflows for operators already familiar with the Darktrace Threat Visualizer. The SaaS console provides access to global maps of activity, detailed logs of user activity and visualization of anomalous chains of behavior.

Each module comes with its own configuration considerations and the implementation process required to ensure operability will vary accordingly. Processes that vary across modules include, but are not limited to: means of data retrieval; permissions requirements; limitation considerations; and license requirements. The specifics configuration requirements of each Darktrace / IDENTITY module and detailed instructions for each setup can be found in the relevant technical documentation in the Customer Portal. Customers should consult this documentation ahead of attempting deployment.

4.2.1.Considerations

Third-party platforms may limit the events in the applicable modules that Darktrace is able to access. Each module has its unique set of considerations that should be reviewed before deploying the module. Delays may be incurred where the external platform does not make events available to the Darktrace module for processing and analysis within the expected timeframe. It is a Customer responsibility to identify and minimize latency between the relevant Darktrace / IDENTITY module and a given third-party platform, Darktrace does not bear responsibility for latency resulting from third-party platforms. Latency between event occurrence and when it was made available to the module are indicated in the event metadata within the Threat Visualizer.

4.2.2.Permissions

Each Darktrace / IDENTITY Module has its unique set of permissions that are required to deploy the module. The permissions vary by third-party platform and details about the permissions needed can be found in the relevant technical documentation in Customer Portal. Different permissions apply in relation to data retrieval than apply in relation to taking a response action.

4.2.3.Deploying Autonomous Mode

Darktrace / IDENTITY modules must be reauthorized to enable Darktrace Autonomous Response functionality after a Darktrace Autonomous Response license key is added, regardless of whether this permission was present at initial authorization. Additional permissions are required in order for Darktrace to take actions within a given third-party platform. The permissions vary per vendor as outlined above,

and the information needed on the Darktrace end can be found in the relevant technical documentation in Customer Portal.

Whether Darktrace / IDENTITY can take action autonomously, or must wait for human approval, is defined by a schedule or model setting. The seven-day, hourly timetable allows blocks of autonomous action or enforced human oversight to be scheduled. For some high severity activity types, it is recommended to ensure Darktrace / IDENTITY Autonomous Response can always take an autonomous action regardless of the scheduled state; in this case, an override can be configured in the model setting.

The Customer is responsible for configuring Darktrace / IDENTITY Autonomous Response to run fully autonomously. Fully autonomous mode is the end goal of all Darktrace / IDENTITY Autonomous Response deployments; reaching a fully autonomous state where the AI can take action whenever unusual or concerning behavior is detected, without the need for human oversight. This mode lends itself to a minimal-interaction workflow, where the Customer may infrequently need to modify actions through the Darktrace Threat Visualizer interface, API, or Darktrace Mobile App, but on the whole can leave Darktrace / IDENTITY Autonomous Response to operate with little intervention.

4.2.4. License Requirements

Some third-party platforms require special licenses to allow events to be shared with Darktrace. For best results the license requirements should be reviewed before implementing the Darktrace / IDENTITY module. The requirements are listed in the Darktrace / IDENTITY section in Customer Portal.

4.3. External Alerting, Mobile App and Model Editing

Details on External Alerting, use of the Darktrace Mobile app and managing edits made to Autonomous Response models can be found in the applicable sections of the Darktrace / NETWORK product specification.

5. Administration

Detailed information about individual administrative tasks is outlined in the relevant documentation for the deployment of the specified component, module, or administrative action intended to be performed.

General administration of the Darktrace deployment is performed in the Threat Visualizer interface:

- The majority of configuration including administration of system settings such as proxies, authentication configuration such as LDAP and SAML SSO, deployment of alert and threat intelligence integrations, and other administrative tasks are performed in the “System Config” page.
- User and group management, including the assignment of data visibility and permissions, is performed on the “Permissions Admin” page.
- The application of tags to user entities is performed on the “Tag Manager” or “Profiles” pages.
- System health information and system alert resolution is performed on the “System Status” page.

Darktrace physical and cloud instances are each seeded with random passwords and two-factor authentication secrets at build time. These initial secrets are stored by Darktrace. Darktrace will provide the client with username/password combinations granting access to the Darktrace Threat Visualizer interface and, for physical instances, the SSH administration console (“Darktrace Console”); these passwords can be optionally changed by the client at any time. Access to the underlying backend systems of the Darktrace instance is reserved for Darktrace only.

5.1. Administration of Darktrace / IDENTITY Modules

Darktrace / IDENTITY modules are configured from the Darktrace “System Config” page. Each module has an individual entry with module-specific configuration settings. Users can modify Autonomous Response eligibility, user activity filters, and authenticate - or reauthenticate - the module from this location. It is a Customer responsibility to ensure that each module is configured correctly to ensure correct operation and interactivity with the associated third-party platform.

5.2. De-Commissioning

The process for removing a Darktrace / IDENTITY module will differ between third-party platforms.

Modules can be “de-authenticated” from the Darktrace System Config page - this will remove any authentication information used by the Darktrace ActiveAI Security Platform to contact the third-party platform. After de-authentication, any components created during the authentication process such as OAuth app registrations, service principals, API keys, or other authentication methods can be removed or deleted by the client from the relevant third-party platform. Darktrace strongly recommends that the client remove any remaining components when Darktrace / IDENTITY module service is ended.

6. Requirements

For specific technical requirements for a given module, consult the Customer Portal for the relevant technical documentation.

7. Customer Responsibilities

Notwithstanding any additional responsibilities set out in any other section of this Product Specification, and without limitation, Customer’s responsibilities are to ensure that:

- The Darktrace / IDENTITY module has been configured and authorized according to Darktrace instructions.

This includes ensuring that all permissions requested by the module are granted, whether during an authentication flow such as OAuth, granted separately as permissions or roles in the third-party platform, or granted/ provided in another comparable deployment process.

- Authentication is performed by - or granted via - a user with the appropriate permissions as outlined in the relevant Darktrace / IDENTITY module documentation.

Recommendations or requirements placed on authorizing users are due to technical necessities. Darktrace cannot guarantee that OAuth grants, API keys, or other authentication methods performed by users without the appropriate permissions will result in functional module operation.

- The Darktrace / IDENTITY module retains authentication and, if lost, is re-authenticated in a timely manner. Unauthorized modules are unable to continue with monitoring or Darktrace Autonomous Response.
- The Darktrace / IDENTITY module is re-authenticated when requested by Darktrace. The module may require re-authentication to add additional capabilities (for example, after Autonomous Response is licensed).
- The Darktrace / IDENTITY module operates within the maximum threshold for event and user coverage, and that deployment components are not overloaded.

- Modifications are not made to the third-party platform configuration or licensing that would prevent operation.

Examples include the removal of auditing within the platform, removal of licenses required to contact specific APIs, removal of permissions from associated API keys, or any other configuration alteration that would interrupt service operation.

- System health issues are monitored using the provided tools, then if alerts arise, addressed and rectified in a timely manner.
- End-user access is managed appropriately and within recommended best practice.
- Changes made to components or overall system configuration by users of the Darktrace platform do not impact the system's ability to function, or otherwise degrade service.

Examples of user error which may result in degraded service include the exclusion of key users using activity filters, the exclusion of all users from Darktrace Autonomous Response, the creation of excessively overactive models, and the modification of existing models to the extent that alerts criteria can no longer be met.

- Darktrace / IDENTITY Autonomous Response has been deployed in line with Darktrace best practice recommendations and, if changes are made to the third-party platform configuration, ensure the deployment continues to meet best practice recommendations across its lifecycle.
- Darktrace / IDENTITY Autonomous Response is enabled globally, in the deployment settings, and in any other applicable locations to ensure actions can be taken.

Deployment configuration must also not preclude the taking of reasonable actions, for example, where all users are placed in an "immune" state.

- Darktrace / IDENTITY Autonomous Response is granted necessary autonomy to enact actions, or, where human confirmation mode is partially or fully in place, ensure that actions are activated by human operators in a timely manner.
- Alterations are not made to IDENTITY Autonomous Response models or models utilized by IDENTITY Autonomous Response that would significantly impede the ability to create targeted actions.
- IDENTITY Autonomous Response actions are not created excessively by custom models or existing models modified to contain actions.

8. Considerations

Due to the variation in third-party platforms, considerations will differ between individual Darktrace / IDENTITY modules. Please refer to the individual module documentation for further information.

Delays may be incurred where the external platform does not make events available to the Darktrace module for processing and analysis within the expected time frame. Delays of this nature are the responsibility of the associated third-party platform vendor. Latency between event occurrence and when it was made available to the module are indicated in the event metadata within the Threat Visualizer.

Where delays are incurred receiving events from the third-party platform, Darktrace / IDENTITY modules may be unable to take Autonomous Response actions in a timely manner due to the latency between event creation and its ingestion and processing by Darktrace / IDENTITY.

If latency exists between Darktrace / IDENTITY and the third-party platform (for example, due to limited outbound bandwidth or delays in the third-party platform processing of requests), Darktrace / IDENTITY Autonomous Response actions may incur delays.

The responsibility to ensure that Darktrace / IDENTITY has an appropriate level of outbound access and bandwidth lies with the Customer. The responsibility for any delays incurred between instruction by Darktrace / IDENTITY and action occurrence in the third-party platform lies with the associated third-party platform vendor.

Individual vendors may place rate-limiting restrictions on the APIs and methods utilized by Darktrace / IDENTITY modules - please refer to the individual module documentation for further information.

Darktrace / IDENTITY modules monitor and retrieve data from third-party platforms using the APIs and integration methods presently offered by those platforms. If this coverage or available capabilities are modified or revoked by the third-party, Darktrace will endeavor to maintain continuity but is ultimately limited in capacity to restore service.

Darktrace Autonomous Response is subject to the restrictions placed by third-party platforms on API capabilities. This can result in specific user types being ineligible for actions, specific circumstances impacting the efficacy of actions, and the requirement for additional configuration (in both Darktrace and/or the third-party platform). Individual considerations are outlined in the individual module documentation.

Changes made to third-party platform capabilities or APIs are outside the scope of Darktrace control and may be introduced at any time. Changes to the known considerations due to alterations by vendors of these platforms may arise at any time.

In many third-party platforms, Darktrace / IDENTITY modules are reliant upon specific vendor licenses to retrieve the events necessary for operation or to take Autonomous Responses.

If the required licenses are not present, or are revoked, degradation or interruption of service will occur.

Darktrace / EMAIL

1.	Product Features	71
1.1.	High Level Summary	71
1.2.	Data Retrieval and Ingestion	72
1.3.	Darktrace Analysis	73
1.4.	Output	77
2.	Deployment Architecture	84
3.	Implementation	84
3.1.	Initial Setup process by Email Provider	84
3.2.	Implementing Autonomous Actions	85
3.3.	External Alerting, Mobile App and Model Editing	85
4.	Administration	85
5.	Requirements	86
5.1.	Requirements of Darktrace /Email for Microsoft 365	86
5.2.	Requirements of Darktrace / EMAIL for Hybrid Exchange	88
5.3.	Requirements of Darktrace / EMAIL for OnPremises Exchange	88
5.4.	Requirements of Darktrace / EMAIL for Gmail	89
6.	Customer Responsibilities	89
7.	Roles	91
8.	Considerations	91

1. Product Features**1.1. High Level Summary**

Darktrace / EMAIL analyzes email communications for anomalous, potentially malicious content. On detection the system will make an evaluative decision about the best response. Darktrace / EMAIL will then instruct the email third-party provider to enact the decision through methods made available by that third-party provider. The decisions will vary according to various detection parameters but may include removal of the email message from the inbox, movement of the email message within the inbox, automatic neutralization of email content such as attachments, and end-user notification messages. Darktrace / EMAIL also provides Account Protection which, in supported third-party systems, will report on anomalous usage of user accounts.

1.2. Data Retrieval and Ingestion

Darktrace/EMAIL does not operate as a “gateway” or “inline” of email traffic; third-party email providers send copies of email to Darktrace /EMAIL for analysis and, under relevant circumstances, mitigation.

Darktrace/EMAIL ingests original copies of each inbound email received and outbound email sent by an organization that is within the scope of its monitoring. Copies are received using supported methods provided by the third-party email vendor; the exact methodology of configuration will vary according to the specific third-party and customer preference.

Unless explicitly excluded, Darktrace/EMAIL receives all email traffic including inbound email (originating outside the organization), outbound email (leaving the organization), and lateral email (internal to internal).

In all operating modes, Darktrace/EMAIL requires access to third-party APIs to retrieve contextual data, historic mailflow data, and other information relevant to its analysis. This access is provided through mechanisms appropriate to the specific platform and can include OAuth application registration, service user creation, and other methods.

1.2.1. Journalling

Darktrace/EMAIL can utilize ‘journalling’ to automatically receive copies of emails in transit to/from the organization. Journalling is a native functionality provided by vendors such as Google Gmail and Microsoft Exchange on-premise for archiving purposes.

In order to receive email traffic, Darktrace will provide a unique email address (“journalling address”) to act as the recipient of forwarded mailflow from the organization.

Journalling is used in combination with mailflow rules and other routing components defined during deployment to ensure emails are routed to Darktrace/EMAIL successfully. The nature of these routing components will differ between email providers.

1.2.2. API Ingestion (restricted availability)

In specific platforms, Darktrace/EMAIL can instead ingest email traffic via APIs provided by the third-party platform when granted the appropriate permissions. This operation mode may also be referred to as “Quick Setup Mode”.

In this mode, Darktrace/EMAIL is notified by the third-party platform when new email messages are created, modified, or sent (“change notifications”). Any delays in the creation of this notification are the responsibility of the third-party platform.

1.2.3. Limiting Darktrace/EMAIL Scope

Emails sent to Darktrace /EMAIL may be restricted by configuration within the third-party email provider, or by configuration within the Darktrace/EMAIL interface.

- In journalling mode, restrictions on the emails routed to Darktrace/EMAIL are built into the mailflow components created during the deployment process.
- In API ingestion mode, restrictions can be made on the Darktrace/EMAIL interface “Config” page. Darktrace/EMAIL will not retrieve email data via API for mailboxes outside the scope of this configuration.

Unless explicitly excluded using a method outlined above, Darktrace/EMAIL receives all email traffic including inbound email (originating outside the organization), outbound email (leaving the organization), and lateral email (internal to internal).

As of Darktrace/EMAIL April 2024, lateral email traffic may be processed and displayed alongside inbound and outbound mailflow in the Darktrace/EMAIL interface.

Permissions are provided to restrict visibility to lateral email messages in the Darktrace/EMAIL interface. Specific users and groups can also be opted out from the Darktrace/EMAIL interface “Config” page.

1.2.4. Historic Email/Contextual Data

On initial deployment, Darktrace/EMAIL will also retrieve historic email data for retrospective processing. Darktrace/EMAIL retrospectively processes email metadata to gain an intimate knowledge of users, email addresses, correspondents, and routine operations. This processing allows Darktrace/EMAIL to create historic behavioral profiles for use in analysis before suitable quantities of mailflow have been observed.

Retrospective data is retrieved using API access granted to Darktrace/EMAIL during the configuration process; the mechanism by which this API access is granted differs between third-platforms.

Contextual data such as user aliases, personal names, and user groups will be periodically retrieved after initial deployment where relevant to the operation of Darktrace/EMAIL analysis or relevant to investigation workflows in the Darktrace/EMAIL interface.

1.2.5. Account Protection

In supported environments, Darktrace / EMAIL can also ingest activity of the user accounts associated with the customer email domains. Activity data is retrieved via APIs or export formats provided by the third-party platform.

This feature - referred to as “Account Protection” - operates a limited subset of the capabilities outlined under Darktrace / IDENTITY.

1.3. Darktrace Analysis

Darktrace/EMAIL inspects email at the point it transits through the organizational mailflow environment. It extracts hundreds of data points from both the raw email content and historical email behavior of the sender and the recipient. These metrics are combined with ‘pattern of life’ data of the intended recipient, or sender, sourced from other Darktrace Real-Time Detection coverage areas.

The combined set of metrics are passed through machine learning algorithms to produce a single anomaly score of the email, and various combinations of metrics will attempt to fire model alerts which will help define the ‘type’ of email. In conjunction with the specific model alerts and the Darktrace/EMAIL Anomaly score, the system may provoke actions upon the email designed to prevent delivery of the email or to neutralize potentially malicious content. Campaign detection - where Darktrace/EMAIL identifies a flood of emails with similar malicious intent - can also influence these actions retrospectively to ensure full containment.

Darktrace analysis is multi-stage. The following outline highlights key elements of the process but is not intended to be an exhaustive list of all analysis performed during operation of the Darktrace platform, Threat Visualizer interface, or any underlying components.

Darktrace / EMAIL will extract metrics from emails observed and combine these with any previously learned behavioral data regarding the historic patterns of elements of that email. This data is passed through AI modelling engines to produce machine learning and statistical outputs. The combined set of machine-learned output and raw data is evaluated to produce an 'anomaly score'.

This anomaly score, the raw email data, historical behavioral patterns and the machine-learned outputs are finally evaluated to produce a suggested action which can then be applied as an instruction to the third-party email provider. The analysis will produce a suggested action proportional to the perceived likelihood of threat and accounting for the potential for false positives.

1.3.1. Darktrace / EMAIL Metrics

Darktrace/EMAIL extracts and computes over one thousand metrics for ingested email messages; each metric is continuously calculated and assessed as messages are processed. Metrics are computed through a combination of statistical analysis, textual analysis, and both supervised and unsupervised machine learning-based approaches (standalone and multistage).

Metrics are features Darktrace/EMAIL has derived for both the email message (including associated content) and the wider context of the interaction. Metrics may, for example, focus on the sender, the recipient, historic behavioral patterns between the sender/recipient domains, the email transit path, the composition of attachments, or the structure of any links and the appearance of their end-location. Metrics can include raw extracted data from the message (e.g., a sender domain), an output from one a machine learning classifier (e.g., a score of how closely the recipient is associated with the sender) or a statistically derived identity metric that determine how well known the sender or communicator is to the organization.

Metrics may be reliant on, or contribute to, overall 'pattern of life' calculations.

Darktrace provides the ability to utilize the output of this analysis using the *models* framework.

1.3.2. 'Pattern of Life'

Darktrace will create individualized behavioral profiles for the user entities it observes and surface activity which is considered inconsistent with the expected norm. Darktrace/EMAIL continuously computes this 'pattern of life' at multiple levels, including (but not limited to) sender, recipient, peer group, organization, and domain-level.

This baseline is derived from - but not limited to - an analysis of the behavior of the individual entity, analysis of one or more clusters created from similarly behaving entities, and many variable factors such as time. The 'pattern of life' data is continually updated in real time, and reflects the data that has been received, with a greater weighting to more recent data.

Analysis is performed in the "classifier" stage by a multitude of classifiers. This core analysis applies many approaches including unsupervised Machine Learning techniques such as Bayesian Meta-classification, techniques derived from graph theory and network analysis such as Node/Graph Centrality, approaches derived from statistical analysis such as Spectral Clustering and many other techniques. The previous examples of utilized techniques provided are for illustrative purposes and should not be considered exhaustive.

Darktrace does not offer any capability to access underlying behavioral models or direct classification output. Behavioral profiles associated with 'pattern of life' are used in the computation of metrics and overall "anomaly". The output of some 'pattern of life' models may be available in a normalized format in metrics, where applicable.

1.3.3. Darktrace/EMAIL (Antigena) Anomaly Score

Darktrace/EMAIL computes an overall anomaly score (*formerly* “Antigena Anomaly”) for each email it processes. This anomaly score is computed using a multitude of features including relevant metrics, outputs from Darktrace ‘pattern of life’ modeling, and key detections from the Darktrace/EMAIL Models Framework.

1.3.4. Darktrace/EMAIL Models Framework

A model is used to define a set of conditions which, when met, will alert the system to the occurrence of a particular characteristic or anomalous feature. Darktrace/EMAIL operates a model engine framework which is functionally equivalent to the primary Darktrace ActiveAI Security platform model engine.

Darktrace/EMAIL models are comprised of complex metrics, derived and calculated by Darktrace/EMAIL, represented visually in the model editor as logical flows. Output from the complex anomaly framework is available in accessible, building block format and can be combined with simple conditions and logical expressions to create tailored activity detection. Darktrace/EMAIL models use nested logic (“Recipes”) and can be targeted at specific components of the/EMAIL or generically at the message itself.

1.3.4.1. Darktrace-Maintained Models

Darktrace/EMAIL provides a curated set of default models and recipes as standard, designed and constantly tuned by the specialized Darktrace/EMAIL analyst team. These models are supplied by Darktrace and can be developed and modified by the user in the model editor.

Darktrace-maintained models are continuously updated in parallel with the evolving threat landscape. Darktrace/EMAIL also maintains profiling models which are used to better understand the type of email, the route the email has taken, or the user behind the communication itself. These indicators are used for contextualization purposes and are not directly indicative of a behavioral anomaly.

1.3.4.2. Custom Models

Custom models can be defined by operators of Darktrace/EMAIL to meet any criteria; these are created and maintained in the Darktrace / EMAIL “Models” interface. Care should be taken when defining custom models to ensure that the number of actions produced, tags applied, or other triggered behavior is not excessive and does not impact the system’s ability to function, or otherwise degrade service. Similarly, modification of existing default models should not result in excessive activity or alter the logic to the extent that the criteria can no longer be met. The responsibility to ensure that models created or edited remain within these reasonable expected boundaries lies with the operator.

1.3.4.3. Tags

Tags are utilized by Darktrace/EMAIL to summarize model findings and email features into brief, recognizable categories for quick investigation. Darktrace/EMAIL models apply tags as a standard action.

1.3.5. AI Analyst

Darktrace Cyber AI Analyst does not perform additional analysis on Darktrace/EMAIL model alerts.

The Darktrace/EMAIL Cyber AI Analyst cooperates with the Darktrace ActiveAI Security platform where relevant - see “Cross-Capability / Cross-Coverage Area Analysis”.

1.3.6. Cross-Capability / Cross-Coverage Area Analysis

Where other Darktrace Real-Time Detection coverage areas are deployed, cross-platform analysis is performed by multiple components. Links are created for analysis purposes between network entities and entities modeled by other Real-Time Detection components such as, for example, user entities created by Darktrace / IDENTITY module monitoring known to be associated with a given network entity, or credential entries observed by Darktrace / NETWORK monitoring also observed as part of Darktrace / ENDPOINT coverage.

Darktrace AI Analyst will also link together entities across different coverage areas in the creation of AI Analyst Incidents and AI Analyst Incidents Events. Darktrace AI Analyst may also retrieve additional contextual data during investigation from other components if deployed, such as the retrieval of associated emails from Darktrace / EMAIL.

The output from Darktrace components may be fed to Darktrace Real-Time Detection as part of supplementary Threat Intelligence; Darktrace / Attack Surface Management can provide output of malicious asset identification into the “Watched Domains” list of Darktrace Real-Time Detection to create Model Breach alerts when a network entity accesses an Attack Surface Management-identified domain. The above provides illustrative examples of collaborative scenarios between platform components for reference but is non-exhaustive.

1.3.7. Campaign Detection

Campaigns are groups of malicious communications sharing common factors such as attack vector, message composition, originating sender, or other correlations. Darktrace/EMAIL identifies campaigns using a complex, multi-stage analysis process.

1.3.8. Lateral Email Processing

Darktrace/EMAIL automatically processes - passes for multistage analysis and display in the user interface - the inbound and outbound emails it receives.

Darktrace/EMAIL will also process lateral email traffic where email messages originated outside the organization and were forwarded internally, or where a minimum anomaly criteria are met. To compute this criteria, lateral messages are subject to a limited amount of analysis but will not be subject to full processing unless deemed suitably anomalous.

Users with appropriate permissions may also nominate lateral emails for full processing / subsequent analysis from the Darktrace/EMAIL interface.

1.3.9. Link Analysis

Darktrace/EMAIL analyzes the links and domains observed in every message against the context of business communications, known network activity, and global domain intelligence. These contextualized findings then inform the actions Darktrace/EMAIL takes.

When a user requests access to a link which has been diverted by Darktrace/EMAIL (“locked”), Darktrace/EMAIL will perform additional checks to confirm the end location is not malicious before

permitting access. This deeper analysis can also be triggered manually by an operator, when an end-user reports an email containing a link directly to Darktrace/EMAIL using a supported mechanism.

These additional checks are more intensive and use specialized AI tools to detect hallmarks of malicious endpoints, such as spoofed login pages. During the analysis, Darktrace will access the link and take snapshots of the resulting website for image recognition and analysis.

1.3.10. Message Content

Darktrace/EMAIL does not provide mechanisms to identify specific text or phrases in message body content.

1.3.11. Learning Exceptions

“Learning Exceptions” alter the system reaction to a specific sender and are manually created by an operator.

Learning Exceptions inform Darktrace/EMAIL to no longer perform applicable actions on emails from a combination of sender domain/from address and analysis criteria. In specific circumstances, such as where an email exhibits a significantly higher anomaly score than when the exception was created, Darktrace/EMAIL will ignore learning exceptions to ensure the email is contained. However, care should still be taken to ensure that learning exceptions are not created excessively, or prevent Darktrace/EMAIL from operating as intended.

Any learning exceptions created are the responsibility of the customer operator.

1.3.12. Lists

Lists are sets of entries which can be used within Darktrace/EMAIL model logic or to influence system behavior. For example, lists can be used to inform Darktrace/EMAIL that an existing relationship exists between a domain and an external organization, even if it has not yet been observed or observed only infrequently.

Lists can be defined as containing a specific type of data, such as email addresses or domains, or may be generic and contain any string. Lists can be populated manually or by an automatic system action.

Care should be taken to populate lists appropriately. Any modifications to the content of lists which impacts Darktrace/EMAIL processing or analysis are the responsibility of the customer operator.

1.4. Output

The primary output of Darktrace/EMAIL are automated actions against email traffic.

Additional Darktrace/EMAIL outputs can include email metric data, user activity alerts (“Account Protection”), contextual data (both domain-level and user-level), end-user notifications. Analyzed email activity is surfaced in the Darktrace/EMAIL interface, and is optionally available through the Darktrace Mobile App, compatible output formats, or the Darktrace/EMAIL extension of the Threat Visualizer REST API.

The platform may also output alerts about overall system health and component health.

1.4.1. Darktrace/EMAIL Actions

According to the findings of the previous analysis stages, Darktrace / EMAIL may instruct the third-party email provider to perform an action against an email. The actions are with respect to the severity of the anomaly, perceived threat of the email and/or according to custom or predefined business logic set or enabled by the operator.

- Emails that are perceived to exhibit a high likelihood of threat will be subject to a hold action and Darktrace / EMAIL will then instruct the third-party email provider to remove the email from the third-party system.
- Emails that are deduced to possibly contain a threat, but might have negative impact to the business if it were held, will have targeted actions placed on it, such as the removal or neutralization of specific content or end user warnings.
- Emails which are undesirable from a productivity perspective - yet contain no perceived malicious intent - may be moved and reorganized into user accessible email folders on supported systems.

It is recommended that the operator reviews the actions to ensure they appropriately scoped and are in keeping with the desired outcome.

Actions on email content or delivery location are taken natively within the inbox using third-party APIs. Darktrace/EMAIL cannot take action against outbound emails in default operating mode. Specific limitations may also apply on the type of message and recipient inbox.

Operators of Darktrace / EMAIL have the ability to override the suggested actions and adjust any logical conditions of the evaluations according to their requirements. Editing or changing parameters in the system may impact the evaluation and response and it is the responsibility of the operator to ensure that any configuration does not impede the desired functionality.

1.4.1.1. Example Actions

Darktrace/EMAIL may take one or more proactive or reactive “actions” against email messages. Actions can be triggered by models, action flows, overall anomaly score, campaign detection, or other relevant components.

Actions take many forms, for example:

- Altering the email that met the model criteria in the desired way, such as add a banner warning about malicious intent.
- Taking information from the email or message and store it in lists for later use in modeling.
- Adding an informational tag to the instant message for future investigation context.
- Moving an inbound email to a specific folder in the end-user's inbox.
- Replacing an attachment with a “flattened”, converted format.
- Swapping a link end-location with a URL hosted by the Darktrace / EMAIL infrastructure.

Some actions modify the message itself; others modify where - *or if* - it ends up in the user's inbox.

Delivery actions are actions performed on the entire message at the point of delivery, such as redirecting the email to another user, holding the message before delivery, or directing it to the junk folder.

1.4.1.2. Model Actions

Darktrace / EMAIL models have actions directly associated with them. When the model criteria are met the action is triggered, unless prevented by one of the conditions outlined below in “Downgrading or Negating Actions”.

1.4.1.3. Action Flows

Action Flows can perform specific actions - or prevent certain actions - based on simple action or tag logic. Action Flows are configured on a per-group basis.

1.4.1.4. Action Autonomy and Eligibility

Darktrace/EMAIL offers extensive granularity when enabling automatic actions - actions can be limited to certain users, restricted on specific models, or globally suppressed over a certain level of severity.

Care should be taken to avoid excluding users or groups from Darktrace/EMAIL action eligibility. Failure to enable actions globally, to define appropriate action severity levels, or to suitably include users within the scope of monitoring is the responsibility of the Customer.

1.4.1.4.1. Global Action Status

Darktrace / EMAIL will only take actions against emails when actions have been enabled globally. Once enabled, per-recipient eligibility will then be assessed before actions are taken.

1.4.1.4.2. Action Eligibility

Eligibility can be configured at the *group* level - actions are applied to *eligible* recipients. Groups allow for both granular inclusion and exclusion from actions - Darktrace/EMAIL will only take action against recipients where eligible.

Darktrace/EMAIL retrieves user groups from the email provider and allows for the creation of new, custom groups. These smaller group units can be used to change settings or configure actions on a per-department or per-tenant basis.

Other mechanisms to alter, downgrade, or prevent actions entirely are outlined below.

1.4.1.4.3. Downgrading or Negating Actions

Darktrace provides many mechanisms to alter the actions taken by Darktrace/EMAIL, ranging in granularity and complexity:

- “Counter actions” - actions which explicitly prevent other actions of their specific type - can be used to negate specific actions when model or action flow criteria are met.
- Action severity controls the intensity of actions Darktrace/EMAIL can perform in each category; limiting action severity will downgrade actions that are not permitted by the level set. Action severity can be defined on a global, group, or user level.

Modifying the actions applied by default Darktrace/EMAIL models directly within the model is possible but is not advised. The responsibility for selecting an appropriate alternative action, and for any actions created as a result of a custom model configuration, lies with the operator.

1.4.1.5. Retrospective Actions

Darktrace / EMAIL evaluates emails as they are received from the third-party provider. If subsequent evidence is seen which may have altered the evaluation of a previous email, then Darktrace / EMAIL will take retrospective action upon emails present in inboxes.

1.4.1.6. Action Methodology

Actions are taken via API calls to the third-party platform. API access is configured and authenticated during the setup process to provide Darktrace / EMAIL with appropriate levels of permission (in the third-party platform) to perform the required APIs calls.

If changes are made to the components created during configuration - for example, revoking of roles or removal of the authorizing user - necessary permissions and access must be restored to Darktrace / EMAIL to ensure actions can continue to be taken.

Darktrace / EMAIL can only take action against emails received by supported user/mailbox types (*provider-dependent*).

Any failure to apply actions due to upstream delays or failures in third-party APIs is the responsibility of the third-party platform.

1.4.1.7. Manual Actions

Darktrace / EMAIL provides mechanisms to create manual “hold” actions on a per-email or bulk basis. These actions can be created by a user with appropriate permissions from the Darktrace / EMAIL interface or from the Darktrace Mobile App.

1.4.1.8. “Release” and Release Requests

The “release” mechanism allows an operator to release an original, unactioned copy of an email for delivery. Emails may be released from the Darktrace / EMAIL interface or from the Darktrace Mobile App.

Emails can be released to the original recipient or a nominated email address. Darktrace/EMAIL provides user permissions to restrict the ability to release emails to other recipients, or release entirely.

“Release” requires the raw email content to be present in storage and not outside retention limits.

Darktrace/EMAIL will modify message elements such as links and attachments and may remove emails entirely from end-user inboxes which exhibit suspicious properties.

Release requests allow users to request an original, unedited copy of an email message. Darktrace/EMAIL operators can then review the request, perform any necessary investigations, and finally approve or reject it. The option for end-users to directly release held emails is also available.

Operators should ensure that only appropriate email content which does not display malicious or otherwise concerning properties is released to end users. The responsibility for email content released through any of supported mechanisms lies with the approving / initiating user.

1.4.2. Alerts

Darktrace/EMAIL does not create alerts when model criteria are met. Models and model alerts are instead used as part of the analysis process. Darktrace/EMAIL model alerts can be viewed within the email interface on a per-email basis.

Operators may add “notification” actions to models, or configure action flows to generate notification actions, based upon specific criteria. Refer to “End-User Communications” below.

1.4.2.1. Account Protection Alerts

Account protection alerts are created for high-severity models which detect anomalous or concerning user behavior. These alerts can be consumed in the relevant dashboard or user profile.

Alerts surfaced in Darktrace/EMAIL are only those relevant to Darktrace/EMAIL or that indicate anomalous account login activity.

1.4.2.2. External Alerts

Darktrace/EMAIL can issue two primary alert types to external systems - Darktrace/EMAIL alerts and Account Protection alerts.

Account Protection alerts are equivalent to standard model alerts issued for Darktrace/IDENTITY through the mechanisms outlined for the Darktrace ActiveAI Security platform.

Darktrace/EMAIL alerts are generated for all email - both inbound and outbound - observed by Darktrace/EMAIL. Each alert contains a number of key metrics such as sender & recipient(s), Darktrace/EMAIL anomaly scores, successfully applied actions, and details of any links and attachments.

A secondary alert type - Darktrace System Status alerts - may also be generated where the Darktrace platform is experiencing degraded service due to health issues, invalid input, or where it is necessary to highlight changes to system administrators. Alerts include details of the originating host, the severity of the event, and relevant links to investigate or resolve the issue. Notifications are sent when a system event becomes active and can optionally be sent on resolution.

1.4.2.3. High-Severity Model Alerts

Where a full Darktrace/IDENTITY module of a compatible type is also monitoring the users observed by Darktrace/EMAIL, the module will populate high severity model alerts from Darktrace/EMAIL into the user activity logs for relevant users. This data is available in the Darktrace / IDENTITY console and other relevant locations.

1.4.3. Darktrace / EMAIL Cyber AI Analyst

For each email message, Darktrace Cyber AI Analyst breaks down Darktrace/EMAIL's analysis into short, textual summaries (a “narrative”) that allow a user to quickly grasp the message context and any potential threat posed. This information is displayed in a number of locations including the Darktrace / EMAIL interface, the Darktrace mobile app, release request pages, used in responses from the Security Mailbox Assistant, and displayed when an email is analyzed with the Inbox Analysis Add-in (see below).

1.4.4. Email Log Data

Darktrace / EMAIL will store a subset of the data used in the data analysis for later retrieval, search and analysis. Darktrace/EMAIL provides comprehensive access to every email it observes through the Darktrace/EMAIL interface.

Each entry displays high-level outputs from Darktrace/EMAIL's analysis and an overall anomaly score.

Email log data is highly filterable, supporting both simple, quick-apply filters and complex logical queries utilizing over 250 metrics. Filtering the Email Logs also creates a visualization of organizational email traffic over the specified period.

A detailed email information view is also available where operators can review outputs from Darktrace/EMAIL's analysis, approve or reject release requests, view if the email is part of a campaign, and manually change the actions taken.

1.4.5. Mailbox Inspector

The Mailbox Inspector feature provides a live view of the email messages present in an end-users mailbox from the Darktrace/EMAIL interface. This view is permission-restricted.

The data displayed in this view is limited to the data made available to Darktrace/EMAIL by third-party APIs.

1.4.6. Inbox Analysis Add-in

The Darktrace/EMAIL Inbox Analysis Add-in is a Microsoft Outlook Add-in - it does not require user access to the Darktrace / EMAIL interface to operate. End-user interactions with the Inbox Analysis Add-in are surfaced in the Darktrace / EMAIL interface.

The Inbox Analysis Add-in allows end-users to invoke Darktrace/EMAIL analysis from the inbox. The capabilities of the Add-in are highly customizable - users may also request the release of original emails and report emails to the security team for review.

The Inbox Analysis Add-in is only available for specific deployment types and will only operate in supported Microsoft Outlook clients. The Darktrace/EMAIL Inbox Analysis Add-In does not support encrypted or archived email messages.

1.4.7. Misdirected Email Add-in

The Darktrace / EMAIL Misdirected External Add-in is a Microsoft Outlook Add-in - it does not require user access to the Darktrace / EMAIL interface to operate. End-user interactions with the Misdirected Email Add-in are surfaced in the Darktrace / EMAIL interface.

The Misdirected Email Add-in is triggered every time an email is sent. If the system detects that the intended recipients may be incorrect when 'pattern of life' is taken into account, the Add-in will prompt the user to correct them. The Misdirected External Add-in will also provide an alternative recipient address where possible.

The Misdirected Email Add-in is only available for specific deployment types and will only operate in supported Microsoft Outlook clients.

1.4.8. End-User Communications

Darktrace/EMAIL will send messages to end-users or the security team in a number of circumstances. For example, to notify standard users of a held email or to provide a list of actioned emails. Similarly, security team members may receive notifications when Data Loss models are triggered, or when a user reports an email through the Darktrace/EMAIL Inbox Analysis Add-In.

Darktrace provides mechanisms to alter the appearance and content of many automated communications. These modifications are made in the “Messages” page of the Darktrace / EMAIL interface.

1.4.9. Security Mailbox Assistant

The “Security Mailbox Assistant” provides automated, AI-powered triage and response to end-user reports.

When a user reports an email as suspicious through a supported method, Darktrace/EMAIL will perform deep-dive analysis into the content from the reported email and identify similar emails that might be part of a campaign. Darktrace/EMAIL can then reply directly to users, explaining why the reported email was anomalous or found not to exhibit suspicious properties. Any campaigns discovered from user reports are automatically remediated and surfaced on the campaigns page.

In the default operating mode, Darktrace/EMAIL will respond to emails sent to a specific address. Responses are not limited to those forwarded to this address - the Security Mailbox Assistant can respond to end-user reports through multiple means including the add-in, native email-provider reports, and other compatible email security tools.

1.4.10. Locked Link Pages

Darktrace/EMAIL “lock” and “double lock” actions replace URL destinations in email content with a redirection to a Darktrace/EMAIL-hosted page. Further details about this capability are outlined below. Darktrace provides mechanisms to alter the appearance and content of this page.

Limited customization of page behavior is also provided - any changes to default “lock” or “double lock” behavior is the responsibility of the customer operator.

1.4.11. Raw Email Data

Darktrace/EMAIL stores raw email data in a rolling buffer. Original and processed raw emails may be retrieved via the system (according to operator permissions and retention policy). This data is used for actions which require an original copy, such as releasing a held email or original attachment.

Operators may release raw emails back into the original inbox, release to a specific individual, or download from the Darktrace/EMAIL interface directly. A preview of raw email message or attachment content can also be displayed to end-users in the Darktrace/EMAIL interface.

Permissions are provided to restrict access to raw email data or previewed email data. Options to view or download email data are not shown to Darktrace personnel.

Darktrace also provides mechanisms to prevent raw email traffic from being stored by Darktrace/EMAIL at all (data retention). Any modification to the storage and retention period of raw email traffic will impact Darktrace/EMAIL functionality and is not advised unless explicitly required.

1.4.12. Data Exports

Data exports allow information about emails that meet a specific criteria to be exported for review in .csv format. Exports can be downloaded from the Darktrace/EMAIL interface or sent to specific recipients by email.

1.4.13. “Locked” Link Content

Darktrace/EMAIL “lock” and “double lock” actions replace URL destinations in email content with a redirection to a Darktrace/EMAIL-hosted page.

Any attempt by an end-recipient to access the link is instead directed to a Darktrace/EMAIL-hosted page requiring confirmation before any possible redirection to the final destination. In specific cases, further analysis of the destination may be performed and the user may be prevented from accessing the final destination.

Darktrace /EMAIL infrastructure acts as a proxy to the destination after the user has requested access and records the access attempt. Darktrace /EMAIL does not host content from the final destination of “locked” or “double-locked” links redirected via the service.

2. Deployment Architecture

When deployed against cloud based third-party email providers, Darktrace / EMAIL is provided as a SaaS (software as a service) solution. When deployed to protect on-premise Microsoft Exchange systems, Darktrace / EMAIL is provided on dedicated hardware. Hybrid Microsoft Exchange solutions may be deployed as either cloud or hardware depending on the customer environment.

Darktrace / EMAIL is deployed outside the flow of email. It does not sit inline. Analysis of email is made on copies, either transmitted to, or retrieved by Darktrace / EMAIL through native methods provided by third-party email providers. Actions are performed by utilizing API methods provided by the third-party systems.

3. Implementation

3.1. Initial Setup process by Email Provider

Darktrace provisions Cloud or On-premise infrastructure to support Darktrace / EMAIL as part of the Customers Darktrace ActiveAI Platform. The process for implementation then differs by the Customer's email provided; Requirements for each are detailed in Section 5. For optimal detection and response, Darktrace recommends implementing the connector and journaling deployment method.

The Darktrace Customer Portal contains technical guides for implementing Darktrace / EMAIL according to each available email provider. Customers should follow the best practices set out in these guides when implementing Darktrace / EMAIL.

3.1.1. Darktrace / EMAIL for Microsoft 365

3.1.1.1. Darktrace / EMAIL for M365 Quick Setup:

<https://customerportal.darktrace.com/product-guides/main/office-365-deployment-process>

3.1.1.2. Darktrace / EMAIL for M365 with Connector and Journaling:

<https://customerportal.darktrace.com/product-guides/main/antigena-email-connector>
<https://customerportal.darktrace.com/product-guides/main/antigena-email-journal-rule>

3.1.1.3. Darktrace / EMAIL for Hybrid Exchange Summary:

<https://customerportal.darktrace.com/product-guides/main/deploying-antigena-email-in-hybrid-exchange>

3.1.1.4. Darktrace / EMAIL for On-Prem Exchange Deployment Process:

<https://customerportal.darktrace.com/product-guides/main/onpremises-deployment-process>

3.1.1.5. Darktrace / EMAIL for Google Workspace:

<https://customerportal.darktrace.com/product-guides/main/gmail-deployment-process>

3.2. Implementing Autonomous Actions

After implementation, the desired scope of actions should be applied to the mailboxes covered by Darktrace / EMAIL Customers See section Darktrace/EMAIL Actions for details on how this is configured, and the options Customer has.

Darktrace recommends applying actions with maximum severity levels to the pre-populated “All Users” group - containing all active email users observed by Darktrace/EMAIL - to enable actions at a global scale. Reduction of action severity or exclusions of actions entirely on certain groups or individual users can then be applied as desired.

The Customer is responsible for configuring Darktrace / EMAIL Autonomous Response to run fully autonomously. Fully autonomous mode is the end goal of all Darktrace / EMAIL deployments: reaching a fully autonomous state where the AI can take action on any suspicious or undesired email in scope, without the need for human oversight. This mode lends itself to a minimal-interaction workflow, where Customer may infrequently modify actions through the Darktrace / EMAIL interface, API, or Darktrace Mobile App, but on the whole leave Darktrace / EMAIL to operate with little intervention

3.3. External Alerting, Mobile App and Model Editing

Details on External Alerting, use of the Darktrace Mobile app and managing edits made to Darktrace / EMAIL models can be found in the applicable sections of the Darktrace Customer Portal.

4. Administration

General administration of the Darktrace deployment is performed in the Threat Visualizer interface:

- User and group management, including the assignment of data visibility and permissions, is performed on the “Permissions Admin” page.
- Administration of any Darktrace / IDENTITY modules deployed to provide Account Protection capabilities is performed in the “System Config” page.

- Creation and management of group entities, including configuration for reports, action flows, and language settings is performed in the “Groups” page.
- Management of internal domains observed by the Darktrace / EMAIL instance and global action settings are administered on the “System Config” page.
- System health information and system alert resolution is performed on the “System Status” page.

For physical Darktrace appliances of any role, additional administrative tasks such as interface configuration, setting of host variables, and Call-Home configuration may be performed in the appliance console, accessible over SSH. For cloud hosted instances, these administrative actions are managed by Darktrace operations directly and no access to the console is available.

Darktrace physical and cloud instances are each seeded with random passwords and two-factor authentication secrets at build time. These initial secrets are stored by Darktrace. Darktrace will provide the Customer with username/password combinations granting access to the Darktrace Threat Visualizer interface and, for physical instances, the SSH administration console (“Darktrace Console”); these passwords can be optionally changed by the Customer at any time.

Access to the underlying backend systems of the Darktrace instance is reserved for Darktrace only.

Detailed information about individual administrative tasks is outlined in the relevant documentation for the deployment of the specified component or administrative action intended to be performed.

As part of ongoing administration of Darktrace /EMAIL, Customers can schedule Email Threat Reports, export data if required, set custom messages for email actions, run email campaigns, use filters and flags to aid in investigation.

Darktrace/EMAIL administrators can review and modify important configuration settings for data storage, automatic actions and attachment watermarking in the System Config page. Any modifications that are not aligned with Darktrace’s recommended best practices are responsibility of the Customer.

5. Requirements

For organizations who already possess a Darktrace deployment such as Darktrace/NETWORK or Darktrace/ENDPOINT, this Darktrace master instance must be permitted to contact the Darktrace/EMAIL environment over 443/SSL. Your Darktrace/EMAIL environment will be created in an appropriate cloud region; details of the hostnames associated with Darktrace/EMAIL infrastructure can be found in Darktrace/EMAIL Cloud Instance Hosts. The master instance also requires DNS to resolve these domains. For Darktrace cloud-hosted masters, Darktrace operations will ensure this access is in place.

If your organization is subject to restrictions on where mail data should be held (such as the GDPR Legislation in the European Union), please inform your Darktrace representative or note this restriction when submitting information on the Customer Portal.

5.1. Requirements of Darktrace /Email for Microsoft 365

- During the setup process, Customer will be required to log into Microsoft 365 as a **Global Administrator** to grant permission to a Darktrace application to monitor the email environment. The following permissions are requested by Darktrace/EMAIL for email monitoring, autonomous responses, and account protection features.

API	Permission	Scope	Purpose
Microsoft Graph API	Read directory data	Directory.Read.All	Retrieve information on domains and tenants associated with the Microsoft environment.
Microsoft Graph API	Read user profiles	User.Read.All	Retrieve contextual data about users to aid analysis and for display in the Email Console.
Microsoft Graph API	Read all groups	Group.Read.All	Retrieve contextual data about users to aid analysis, for display in the Darktrace/EMAIL Console, and to allow actions to be controlled based upon group membership.
Microsoft Graph API	Read and write mail in all mailboxes	Mail.ReadWrite	Enact Darktrace/EMAIL actions and perform attack engagements (requires Darktrace/Proactive Exposure Management). In API-only mode, also required to retrieve emails.
Microsoft Graph API	Read all usage reports	Reports.Read.All	Retrieve contextual data about users to aid analysis and for display in the Darktrace/EMAIL Console.
Microsoft Graph API	Read all user mailbox settings	MailboxSettings.Read	Retrieve contextual data about users to aid analysis and for display in the Darktrace/EMAIL Console.
Microsoft Graph API	Read and write calendars in all mailboxes	Calendars.ReadWrite	Enact Darktrace/EMAIL actions.
Microsoft Graph API	Read and write all users' full profiles	User.ReadWrite.All	Enact manual actions as part of the Account Protection capability (Disable User).
Microsoft Graph API	Read all of the organization's threat submissions	ThreatSubmission.Read.All	Monitor end-user submissions for the Security Mailbox Assistant.
Office 365 Exchange Online	Use Exchange Web Services with full access to all mailboxes	N/A	Required to support authentication.
Office 365 Exchange Online	Read and write mail in all mailboxes	N/A	Enact Darktrace/EMAIL actions.

- If Customer is configuring mailflow Journaling, a user who is able to access Microsoft admin environments such as the Exchange Admin and Microsoft 365 admin center is also required.

5.2. Requirements of Darktrace / EMAIL for Hybrid Exchange

- A Microsoft 365 (formerly Office 365) Business Essentials license or above is required to access the necessary configuration settings.
- Darktrace/EMAIL currently supports Exchange environments running Exchange Server 2016 and above.
- Darktrace/EMAIL retrieves group and user information from Azure AD. Where all or a subset of users are not synced to the cloud Azure AD environment, this information may be retrieved from LDAP integration. This functionality is optional but recommended for best user experience.
- For organizations with both on-premises and Microsoft 365 users, Darktrace/EMAIL requires the hostname of the EWS domain to take actions within on-premises mailboxes. The EWS domain must also have a valid SSL certificate. If mail is routed through exchange servers exclusively for journaling purposes but no mailboxes are hosted on-premise, this is not required.

5.3. Requirements of Darktrace / EMAIL for OnPremises Exchange

- The Darktrace master instance associated with Darktrace/EMAIL must have access to an LDAP server configured - this allows Darktrace/EMAIL to process the users and groups associated with your Microsoft Exchange environment.
- The Darktrace master instance connected to the Darktrace/EMAIL appliance must be running the most recent Threat Visualizer software version (6.1.0 and above).
- Darktrace/EMAIL for Exchange (OnPremise) supports environments running Exchange Server 2013 SP1, or Exchange Server 2016 / 2019 with NTLM(v2) configured.
- Organizations with more than one Exchange server will need to repeat sections of the configuration process for each server to ensure all mailflow is journaled to Darktrace/EMAIL for Exchange (OnPremise).
- Existing Darktrace/EMAIL appliances must be upgraded to use HTTPS/443 Call-Home.
- During the setup process Customer will be required to log into the Exchange Server interface with an account with Domain/Organization Administrator permissions.
- Necessary firewall exceptions to permit the network connectivity requirements outlined below.

Component	Protocol/Port	Direction
User Interface access and communication	HTTPS/443	Inbound/Outbound to Darktrace Master
Network Time Protocol	NTP/123	Inbound or Outbound depending on your NTP environment
Mailflow for analysis	SMTP/25	Inbound/Outbound from your Exchange Server(s)

Component	Protocol/Port	Direction
Darktrace/EMAIL Actions	HTTPS/443	Inbound/Outbound to your Exchange Server(s)
DNS querying	DNS/53	Outbound
Call-Home	TLS/443	Outbound (to specific IPs/endpoint)
Outbound HTTPS	HTTPS/443	Outbound (Unrestricted)
Endpoint analysis on HTTP URLs	HTTP/80	Outbound (only when required for analysis)

5.4. Requirements of Darktrace / EMAIL for Gmail

- Darktrace/EMAIL for Google Workspace is only available to organizations with an Enterprise or Education Plus (formerly Enterprise for Education) License (or above) due to Google Workspace (formerly G Suite) restrictions on Third-Party Email Archiving.
- Darktrace/EMAIL receives journaled mail directly from Gmail servers; Customer's organizational Sender Policy Framework (SPF) record must contain therefore *include:_spf.google.com*.
- During the setup process Customer will be required to log into the Google Workspace (G Suite) interface with an account with Super Admin permissions.

6. Customer Responsibilities

Notwithstanding any additional responsibilities set out in any other section of this Product Specification, and without limitation, Customer's responsibilities are to ensure that:

- Darktrace / EMAIL has been deployed in line with Darktrace best practice recommendations and, if changes are made to the email configuration, ensure the deployment continues to meet best practice recommendations across its lifecycle.
- Email provider licenses are supported by Darktrace.
- Connectors and journaling rules required for the installation are created according to Darktrace instructions.
- In journaling mode scenarios, where required by the third-party email provider, a valid "Undeliverable Address" is defined for undeliverable journal reports.
- All firewall requirements are met for successful setup of Darktrace /EMAIL.
- Autonomous action is enabled for all users within your organization.
- Any latency issues that might be observed are resolved with the email provider.
- Darktrace maintains visibility over all email traffic to be monitored.

- Necessary permissions for Darktrace /EMAIL to be able to take actions are granted.
- Maximum action severity is set for best result in actioning email.
- System health issues are monitored using the built-in features of the Darktrace / EMAIL product, and if alerts arise, are addressed and rectified in a timely manner.
- End-user access is managed securely and within recommended best practice.
- Changes made to components or overall system configuration by users of the Darktrace platform do not impact the system's ability to function, or otherwise degrade service.
- Suitable workflows and processes are in place to remediate any cyber-attacks/compromises that are contained by Darktrace / EMAIL Autonomous Response
- Alterations are not made to Darktrace / EMAIL Autonomous Response models, or other Darktrace models utilized by Darktrace / EMAIL Autonomous Response, which would significantly impede the ability to create targeted actions.
- The outputs of the Darktrace / EMAIL are monitored, and its use complies with all applicable laws and regulations.
- The necessary permissions for Account Takeover Protection have been granted to Darktrace

7. Roles

Customer Role	Responsibility
Project Manager	Coordinates Customer resources as necessary. Serves as the point of contact between Customer and Darktrace. Drives communication from the Customer side. Serves as the point of escalation for issue resolution and service-related issues. \
End User	Leverages Product functionality; monitors & actions output. Provides feedback to other Customer & Darktrace roles.
Email/System Administrator	Ensures Darktrace / EMAIL coverage is maintained. Performs regular checks of data inputs. Responds accordingly to System Status alerts to ensure optimal data quality and delivery.
Customer Portal Primary User	Manages customer portal access and contact information for all other customer roles. Ensures that Service Contacts are verified.

Darktrace Role	Responsibility
Customer Success	Oversee customer's experience using Darktrace. Act as customers' sponsor & escalation path internally with Darktrace.
Account Executive	Facilitate commercial arrangements between Customer and Darktrace.
Solution Engineer	Provides technical expertise to support Account Executive / Key Account Director and Customer Success Manager with their customer needs.
Cyber Infrastructure Engineer	Provide technical guidance during installation and configuration processes.
Darktrace Customer Support	Provides Support Services as per Master Customer Agreement.

8. Considerations

In *journaling* mode, Darktrace/EMAIL utilizes a Microsoft Exchange functionality called "Journaling". To use this functionality, an email address must be provided to send undeliverable journal reports to. Darktrace strongly recommends a dedicated mailbox for this purpose, as Darktrace/EMAIL cannot monitor or action emails to the mailbox used for undeliverable mail reports. Darktrace provides an address that can be used for this purpose if a suitable mailbox is not available.

Due to the nature of mail-enabled public folders, Darktrace/EMAIL cannot observe, analyze or action mail to or from these locations.

Darktrace / EMAIL can only take autonomous actions when granted appropriate autonomy. Darktrace Email will be unable to act on user groups not configured for autonomous actions or if global or domain

actions are disabled. Darktrace / EMAIL may also be limited in its response capabilities if the full set of actions are not enabled.

Darktrace / EMAIL actions are taken via the email provider's API. Darktrace is not responsible for delays in actions due to the email provider once the action has been submitted to the API.

Darktrace / EMAIL's visibility is limited by the data sent to it by the email provider, whether via journaling or API. Mail that does not reach the Darktrace Email instance for any reason will not be processed, visible within the solution, nor actioned.

Even if all mailflow is configured to be sent to Darktrace / EMAIL, visibility into domains might still be limited by the config page. A domain in the tenant must be set to the 'enabled' state for Darktrace to process and action mailflow going to email addresses on that domain.

The Darktrace/EMAIL appliance for OnPremise Exchange should be situated within the same logical network area as the Exchange server to minimize latency in the journaled mailflow. Connectivity between the Darktrace appliance and the Darktrace/EMAIL appliance must also be consistent.

If a hybrid deployment is made up of primarily cloud mailboxes and is installed as a cloud deployment, even if Darktrace sees mailflow for mailboxes situated on-prem it will NOT be able to action those emails unless the hybrid installation steps are followed.

Emails filtered by a gateway prior to reaching Microsoft or Google will not be seen by Darktrace. It is not possible for Darktrace to be positioned 'in front of' that gateway or to 'undo' false positive actions of the gateway or the email provider's native filtering.

Emails from 3rd-party sending services that are utilizing an internal domain but lack the necessary validation (SPF/DKIM) will likely be treated by Darktrace as external emails and are subject to detection and actioning. The validation of these emails should be corrected or Darktrace configurations should be modified to change this default behavior.

Darktrace / EMAIL – DMARC

1.	Product Features	93
1.1.	High Level Summary	93
1.2.	Data Retrieval and Ingestion	94
1.3.	Darktrace Analysis	94
1.4.	Output	95
1.5.	Interface	96
1.6.	Reporting	96
2.	Deployment Architecture	96
3.	Implementation	97
3.1.	Customer Self-Service Implementation	97
3.2.	Darktrace-Initiated Implementation	97
4.	Administration	97
5.	Requirements	98
6.	Customer Responsibilities	98
7.	Considerations	99
8.	Roles	99

1. Product Features**1.1. High Level Summary**

Darktrace / EMAIL-DMARC assists organizations with the implementation of the Domain-based Message Authentication, Reporting, and Conformance (DMARC) standard for email security. The product provides a guided deployment process for DMARC and evaluates the implementation of relevant email security measures, such as DomainKeys Identified Mail (DKIM) and Sender Policy Framework (SPF). Darktrace / EMAIL-DMARC analyzes publicly available DNS records associated with these measures for validity and suitability and provides overall health indicators. Darktrace / EMAIL-DMARC also provides recommendations and guidance to assist the client organization in reaching optimal configuration.

Darktrace / EMAIL-DMARC ingests aggregate DMARC reports and forensic DMARC reports sent by third-party mail processors for the email domains managed by the client. Reports are analyzed to profile the third parties (whether legitimate or malicious) observed sending emails on behalf of the client domains, to identify trends in this observed mail flow, to highlight potential impersonation campaigns (indicated by a spike in observed DMARC failures for a domain), and to surface overall relevant statistics.

1.2. Data Retrieval and Ingestion

Darktrace / EMAIL-DMARC retrieves and/or ingests three primary data types: DNS records (SPF, DKIM, and DMARC), “aggregate” DMARC reports, and “forensic” DMARC reports.

1.2.1. Records

Darktrace / EMAIL-DMARC polls publicly available DNS records (SPF, DKIM, and DMARC) for the client’s nominated domains. Records are polled at regular intervals.

1.2.1.1. DMARC Reports

Darktrace / EMAIL-DMARC ingests aggregate reports and forensic reports sent by third-party mail processors for the email domains managed by the client. Reports are only sent by third-party mail processors when indicated in the associated DMARC record for the domain - this configuration must therefore be performed before Darktrace / EMAIL-DMARC can begin ingestion.

Reports are sent by third-party mail processors as part of their voluntary participation in the DMARC protocol. The onus to create and send reports is with the third-party mail processor. Darktrace / EMAIL-DMARC cannot compel a third-party mail processor to generate reports. Darktrace / EMAIL-DMARC is also not responsible for the content of reports generated by third-party mail processors.

1.2.1.1.1. Ingestion Methods

Reports can be ingested via two supported modes: SMTP ingestion and API ingestion.

In SMTP ingestion mode, Darktrace / EMAIL-DMARC will provision an email address unique to the client to receive reports, which is added by the client to the publicly available DMARC record. Third-party mail processors will then send reports directly to this address via SMTP.

In API ingestion mode, a Darktrace / EMAIL-DMARC OAuth application is authenticated by a client administrator. This application grants Darktrace / EMAIL-DMARC permission to retrieve reports sent by third-party mail processors to a mailbox within the client domain via authenticated API requests.

1.2.1.2. Additional Data

Darktrace / EMAIL-DMARC will retrieve additional contextual information about the domains and tenancies associated with the organization. This retrieval is made via APIs and/or export formats provided by the third-party email provider managing the client domains.

1.3. Darktrace Analysis

1.3.1. Analysis of Records

Darktrace / EMAIL-DMARC analyzes record structure and format for compliance with governing standards such as RFC. Record values are also analyzed for validity and suitability, and a risk score is computed (where relevant). This analysis process utilizes statistical, mathematical, and pattern-based analysis techniques. Records are re-evaluated after each poll.

Darktrace / EMAIL-DMARC evaluates records against currently available standards. Standards are maintained by third parties and are subject to change.

1.3.2. Analysis of Aggregate DMARC Reports

Aggregate DMARC reports received from third-party mail processors are analyzed to identify third-party vendors or software, mail sources, and to identify trends or sudden changes in DMARC failure rate. Reports received by a supported method are parsed prior to analysis. This process utilizes statistical, mathematical, and pattern-based analysis techniques.

1.3.3. Analysis of Forensic DMARC Reports

Forensic DMARC reports are ingested but not analyzed at this time.

The content of forensic DMARC reports is defined by the third-party mail processor and may differ. Analysis is therefore limited to the metadata or features included by the third-party mail processor in the forensic report.

1.3.4. Global Domain Analysis

Darktrace / EMAIL-DMARC retrieves global domain information computed by fleet-wide Darktrace / EMAIL analysis for contextual purposes. A Darktrace / EMAIL environment is not required for this information to be accessible.

1.3.5. Cross-Platform Analysis

Darktrace / EMAIL-DMARC identifies third-party vendors or software sending on behalf of the client domain. Identified third parties can be marked as approved or denied.

Where a Darktrace / EMAIL environment is present and monitoring the client domains, Darktrace / EMAIL-DMARC will communicate information about this approval/rejection status for observed third parties. This information is not sent where a Darktrace / EMAIL environment is not present.

1.4. Output

Darktrace / EMAIL-DMARC will provide overall recommendations and guidance to assist the client organization in reaching optimal record configuration.

Outputs from Darktrace / EMAIL-DMARC analysis are surfaced in a dedicated user interface.

1.4.1. Record Analysis

For each domain and record type, Darktrace / EMAIL-DMARC will provide an overall health rating and, where applicable, a risk score (*limited to SPF at present*). Ratings are surfaced in the Darktrace / EMAIL-DMARC interface. Ratings can be optionally sent in an aggregate, summary format via email.

1.4.2. DNS Records or Record Entries

Darktrace / EMAIL-DMARC may provide pre-formatted DNS records or pre-formatted record entries of any type. Darktrace / EMAIL-DMARC may also provide recommendations to modify existing records for validity, improved parsing, or to include/exclude permitted senders.

1.4.3. Third-Party Senders

Third-party senders identified in aggregate reports or permitted by SPF to send on behalf of the organization are surfaced in the Darktrace / EMAIL-DMARC user interface. Operators may mark third-party senders as approved or denied - Darktrace / EMAIL-DMARC will generate recommendations for record modifications in response to this operator feedback.

1.4.4. Mail Statistics

Results from aggregate report analysis are organized by origin, sender, and pass/failure status. Sudden changes to DMARC failure volumes identified by trend analysis will create an alert in the Darktrace / EMAIL-DMARC interface and, optionally, an email alert to chosen recipients.

1.4.5. Alerts

Alerts are optionally generated in summary format for overall domain record health or when sudden changes in DMARC failure rate in aggregate mail flow are observed. Alerts can be restricted to specific domains or specific state changes.

The responsibility to ensure that changes to alert configuration do not prevent the creation of alerts lies with the client operator.

1.5. Interface

The Darktrace / EMAIL-DMARC interface is standalone and accessible via the Darktrace ActiveAI Security Portal.

The interface provides step-by-step guidance on achieving optimal DMARC configuration and includes ongoing evaluation of DMARC, SPF, and other relevant DNS records. Statistics on mail flow, pass/failure status, and third-party senders are also available.

1.6. Reporting

Analysis output and statistics are reported directly on the dashboards of the Darktrace / EMAIL-DMARC interface. Exports such as PDF or CSV are not supported.

1.6.1. Audit Log

Darktrace / EMAIL-DMARC audits changes made by operators in the Darktrace / EMAIL-DMARC user interface. The log is accessible within the interface by administrative users.

2. Deployment Architecture

Darktrace / EMAIL-DMARC is offered in a SaaS architecture model. No on-premises or client-managed deployment architecture is offered.

Infrastructure is provisioned and maintained by Darktrace Operations in a Darktrace-managed public cloud environment. DMARC is hosted in a multi-tenant architecture that utilizes industry-standard practices for data segregation.

Data is encrypted at rest and stored in a region close to the customer in the Darktrace-managed public cloud environment.

There is currently no US Eyes Only Service option for Darktrace / EMAIL-DMARC.

3. Implementation

During any Darktrace / EMAIL-DMARC provisioning, customer will need to confirm the Darktrace-managed public cloud region used to host their instance (Azure US East, EU West, UK South and SE Asia are available). Depending on the Implementation Model, the following differences apply.

3.1. Customer Self-Service Implementation

Darktrace primarily operates a customer self-service implementation model for Darktrace / EMAIL-DMARC, available through the Microsoft Azure Marketplace.

If not already a Darktrace customer, access to the ActiveAI Security Portal will also be provisioned as part of the Marketplace flow.

Once provisioned, a setup guide within the Darktrace / EMAIL - DMARC User Interface serves as the starting point for the configuration and deployment of the product, and for ongoing guidance on how to ensure compliance and best practices across domains.

3.2. Darktrace-Initiated Implementation

In certain limited circumstances, Darktrace may initiate Darktrace / EMAIL - DMARC provisioning by contacting customers directly. Implementation steps depend on whether the customer has an existing Darktrace / EMAIL instance.

3.2.1. Customers with existing Darktrace / EMAIL

If customer has an existing Darktrace / EMAIL instance, either through a contract subscription or Proof of Value engagement, the DMARC tenant-domain information will be pre-populated from the domains set on that EMAIL instance when it was provisioned. These domains are then verified and associated with the tenant ID through a Microsoft API.

3.2.2. Customers without existing Darktrace / EMAIL

If the DMARC instance does not have an associated Darktrace / EMAIL instance, or if the /EMAIL instance does not contain all the required tenant & domain information the customer requires for DMARC, then customer must provide that tenant & domain information to their Darktrace representative, who will update the implementation details accordingly.

4. Administration

Access to Darktrace / EMAIL-DMARC is provided via the Darktrace ActiveAI Security Portal, a consolidated view to access and manage Darktrace products and services.

Administration within Darktrace / EMAIL-DMARC is currently limited to alert recipient configuration, domain group configuration, third party acceptance, audit log visibility and modification of the report ingestion mode.

Darktrace / EMAIL-DMARC contains no direct user administration capabilities. Access and permissions are controlled via the Darktrace ActiveAI Security Portal.

Detailed information about individual administrative tasks is outlined in the relevant documentation for the deployment of the specified component or administrative action intended to be performed.

Access to the underlying systems of Darktrace / EMAIL-DMARC is reserved for Darktrace personnel only.

5. Requirements

Darktrace / EMAIL - DMARC is available exclusively for Microsoft 365 / Office 365 customers. Google Workspace and other email providers are not supported.

Darktrace / EMAIL - DMARC has no dependency on subscriptions to Darktrace / EMAIL or any other Darktrace product or service.

6. Customer Responsibilities

Notwithstanding any additional responsibilities set out in any other section of this Product Specification, and without limitation, Customer's responsibilities are to ensure that:

- Inform Darktrace of cloud region hosting requirements through Azure Marketplace, as well as any changes to requirement to initial Implementation.
- Ensure that a supported method of ingestion is configured and authentication steps are followed fully for all domains.
- In *API Ingestion* mode, ensure that authentication is performed by, or granted via, a user with the appropriate permissions.
- In *API Ingestion* mode, ensure that Darktrace / EMAIL-DMARC retains authentication and, if lost, is re-authenticated in a timely manner.
- In *API Ingestion* mode, ensure that DMARC reports are sent to a mailbox within the domains Darktrace / EMAIL-DMARC is authorized to access.
- Ensure that new domains created after initial authentication are enrolled in Darktrace / EMAIL-DMARC via a supported mode.
- Ensure that any recommended changes to DNS records are made, and that the record content is validated before upload.
- Ensure that changes are not made to DMARC records for monitored domains that prevent or interrupt the DMARC reporting process.
- Ensure that DNS records are managed and maintained appropriately.
- Ensure that only valid third-party vendors are approved to send on behalf of the organization through the Darktrace / EMAIL-DMARC interface.
- Ensure that end-user access to the Darktrace / EMAIL-DMARC interface is managed appropriately and within recommended best practices.

- Ensure that alerting is scoped to appropriate recipients within the client organization.
- Ensure that changes made by operators to components or configuration do not impact the system's ability to function or otherwise degrade service.
- For example, over-stringent or misconfigured alert filters which prevent the creation of any alerts.

7. Considerations

Darktrace / EMAIL-DMARC is unable to provide all functionality if DMARC reports are not being received on behalf of the client domain(s).

Delays may occur from DMARC configuration to the first received aggregate report. Delays are expected to take between 24 and 72 hours.

Aggregate reports are generated on a regular basis and are not expected to be real-time indicators of mail flow.

Forensic report content is controlled by third-party mail processors and may contain email body content, sender information, or other potentially sensitive information.

Participation in DMARC is voluntary and may be implemented differently by different third-party mail processors. Only participating third-party mail processors will produce records. Darktrace / EMAIL-DMARC is unable to compel specific third-party mail processors to generate reports. Darktrace / EMAIL-DMARC is not responsible for how the SPF/DKIM/DMARC records are evaluated by third-party mail processors or other intermediary software. Darktrace / EMAIL-DMARC is not responsible for impacted delivery or sending reputation as evaluated by third-party mail processors or other intermediary software. Darktrace / EMAIL-DMARC is not responsible for the content of DMARC reports generated by third-party mail processors.

Changes made to permitted senders in SPF records should only be made when certain and should only permit valid third-party vendors. Darktrace / EMAIL-DMARC is not responsible for any valid third parties removed from SPF records or invalid third parties added.

Darktrace / EMAIL-DMARC can only evaluate records against currently available standards. Standards may be subject to change.

In *SMTP Ingestion* mode, a Darktrace domain email address must be added to the publicly available DNS record of all participating domains.

Recommendations or requirements placed on authorizing users are due to technical necessities. Darktrace cannot guarantee that OAuth grants, API keys, or other authentication methods performed by users without the appropriate permissions will result in functional operation.

8. Roles

Customer Role	Responsibility
Project Manager	Coordinates Customer resources as necessary. Serves as the point of contact between Customer and Darktrace. Drives communication from the Customer side. Serves as the point of escalation for issue resolution and service-related issues.
End User	Leverages Product functionality, monitors & actions output. Provides feedback to other Customer & Darktrace roles.

DMARC System Administrator	Ensures Darktrace / EMAIL-DMARC coverage is maintained. Performs regular checks of data inputs. Responds accordingly to System Status alerts to ensure optimal data quality and delivery.
Customer Portal Primary User	Manages Darktrace Customer Portal access and contact information for all other customer roles. Ensures that any Service Contacts are verified.

Darktrace Role	Responsibility
Customer Success	Oversee customer's experience using Darktrace. Act as customer's sponsor & escalation path internally with Darktrace.
Account Executive	Facilitate commercial arrangements between Customer and Darktrace.
Solution Engineer	Provides technical expertise to support Account Executive / Key Account Director and Customer Success Manager with their customer needs.
Darktrace Customer Support	Provides Support Services as per Master Services Agreement.

Darktrace / CLOUD

1.	Product Features	102
1.1.	High Level Summary	102
1.2.	Data Retrieval and Ingestion	103
1.3.	Darktrace Analysis	105
1.4.	Output	109
1.5.	Framework Assessment	111
1.6.	Interface	112
1.7.	Usage metrics	112
1.8.	Reporting	113
2.	Deployment Architecture	114
2.1.	Darktrace (ActiveAI Security Platform) Instance	114
2.2.	Hosted Darktrace / CLOUD Infrastructure	114
2.3.	Virtual Sensors	114
3.	Implementation	114
3.1.	Darktrace Cloud for AWS	115
3.2.	Darktrace Cloud for Azure	116
3.3.	Data checks for Darktrace / CLOUD	118
3.4.	Darktrace virtual sensors	118
4.	Administration	120
4.1.	CLOUD Console Administration	120
4.2.	4.2 Darktrace / CLOUD Quick Setup Wizard	121
4.3.	4.3. Administration of Darktrace / CLOUD Modules	121
4.4.	De-Commissioning	121
5.	5. Requirements	121
5.1.	Darktrace / CLOUD Region Availability and Security	122
5.2.	Requirements of Darktrace / CLOUD for AWS	122
5.3.	Requirements of Darktrace / CLOUD for Azure	123

6.	Customer Responsibilities	125
7.	Considerations	128
8.	Roles	130

1. Product Features

1.1. High Level Summary

Darktrace / CLOUD enumerates all resources within a public cloud environment to create a database of enriched cloud assets and construct groupings of interconnected entities called “architectures.” Darktrace / CLOUD analyzes these assets and architectures to detect misconfigured resources, identify critical infrastructure, and monitor the overall health and compliance of the cloud environment.

Darktrace / CLOUD also retrieves and analyzes user behavior in the cloud environment and cloud network activity (either agentless via flow logs or in raw form ingested by Darktrace virtual sensors). This real-time detection evaluates the activity of users and virtual resources in the cloud environment for behavior outside of a “normal” state. This state is created by an ongoing, real-time analysis of connections and other supplied data events across the cloud environment. The system profiles individual entities based on their activity and the activity of those it deems “peers” due to similar behavioral activity (“modeling”).

Alerts are raised when specific system configuration criteria or a minimum threshold of unusual behavior are met. The output is surfaced to operators in the Darktrace Threat Visualizer interface and Darktrace / CLOUD console for investigation and resolution. Darktrace / CLOUD also provides autonomous response capabilities against cloud entities - actions can be taken through native cloud provider APIs or by targeted, connection-based actions.

1.1.1. Terminology

The following concepts will be referred to extensively during the following description.

1.1.1.1. Assets

Assets are discrete entities present in the cloud provider environment that can hold metadata, originate from a supported cloud provider service, and have not been explicitly excluded from enumeration (for example, default entities present in all environments with no relevant use). Assets are discovered by Darktrace / CLOUD during enumeration.

Examples of assets may include virtual machines, virtualized firewall rule policies, access policies, ephemeral resources such as containers, users, and virtual functions or scripts.

Darktrace / CLOUD is continuously expanding the range of supported cloud services and entities across all supported cloud providers.

For the purpose of usage metrics, assets are distinguished by two different types – Identities and Workloads. For further information see the Usage Metrics section.

1.1.1.2. Accounts

“Accounts” is used generically to refer to logically separate groupings within the relevant cloud provider. Equivalence to “account” in AWS and “subscription” in Azure can be presumed.

1.1.1.3. Identities

Users with permissions to act within the cloud provider environment, user-like asset types, and associated assets which impact IAM (Identity and Access Management), are referred to as identities.

1.2. Data Retrieval and Ingestion

Darktrace / CLOUD retrieves cloud resource information, log-based activity data, cost reports, and log-based network activity from cloud providers. This is typically achieved via APIs or export formats provided by the cloud provider. Due to the variation in APIs, authentication methods, and SDKs provided by each platform, the methodology for ingestion for each component will differ between cloud providers. The method of authentication and retrieval is outlined in the relevant documentation provided for Darktrace / CLOUD coverage in each cloud provider environment, as found on the Customer Portal.

Darktrace / CLOUD retrieves raw network traffic for ingestion via Darktrace virtual sensors operating in a supported mode.

1.2.1. Asset Enumeration

Upon initial authentication, Darktrace / CLOUD will enumerate all supported entities within the cloud provider accounts to which it has been granted access. Enumerated entities are recorded as assets in Darktrace / CLOUD.

Enumeration is performed via API access to the cloud provider, granted during initial authentication. Enumeration is performed at regular intervals thereafter. Changes to key cloud services and asset types observed through audit logging will also trigger Darktrace / CLOUD to re-enumerate.

1.2.2. Audit Activity

Data ingested by Darktrace / CLOUD is retrieved directly from the audited activity logs of the relevant cloud provider. Returned information is therefore limited to the events that each vendor chooses to audit and the data recorded as part of those audit log entries. Variation in both the events available in each cloud provider and the level of detail is expected.

Darktrace / CLOUD ingestion also operates in a passive, “open” ingestion mode. Darktrace / CLOUD does not actively request specific events nor impose restrictions on the data retrieved from the cloud provider. The range of events available to Darktrace / CLOUD will therefore vary between each monitored platform and each client. Exceptions exist where the resource types to be monitored are pre-defined during configuration or where individual event types must be removed due to excessive volume.

Latency between event occurrence and Darktrace / CLOUD processing and analysis is impacted by the speed at which the cloud provider makes events available. Delays of this nature are the responsibility of the cloud provider. Latency between event occurrence and when it was made available to Darktrace / CLOUD is indicated in the event metadata within the Threat Visualizer.

1.2.3. Costing

Darktrace / CLOUD may also invoke the creation of, and retrieve, cost reporting generated by the cloud provider. Costing data is retrieved via API access granted during initial authentication. Ingestion of cost reporting is optional.

1.2.4. Cloud Network Traffic Ingestion

Darktrace / CLOUD includes two grades of network coverage: comprehensive traffic monitoring (using Darktrace virtual sensors) and agentless flow log monitoring (using cloud provider mechanisms).

Darktrace 'pattern of life' analysis and Darktrace autonomous response perform most effectively when provided with the depth of metadata generated from raw network traffic monitoring. Darktrace therefore strongly recommends a full traffic mirroring solution which utilizes Darktrace vSensors for comprehensive monitoring overall, or a critical subset of, cloud-hosted networks.

Full visibility requires all cloud-hosted network traffic to be mirrored or ingested by Darktrace in some format. It is the responsibility of the client to ensure that Darktrace maintains visibility over cloud network traffic. The network traffic provided to Darktrace must be of suitable quality and not contain duplication, fragmentary data, or be incomplete in any fashion (for example, unidirectional).

1.2.4.1. Raw Network Traffic Ingestion

Darktrace / CLOUD real-time detection retrieves and processes raw cloud network traffic. This cloud network traffic may be delivered to a virtualized Darktrace vSensor instance via a compatible cloud vendor traffic mirroring offering (for example, Amazon VPC Traffic Mirroring), via host-based osSensor agents, or by another supported mode. Darktrace does not currently provide mechanisms to send traffic directly from a cloud environment to a Darktrace master or Unified View instance.

1.2.4.2. Flow Log Ingestion

Darktrace / CLOUD also supports agentless, pre-processed traffic ingestion in the form of "flow logs" - a high-level, metadata representation of network traffic activity. Flow logs can be generated by the cloud provider on instruction by Darktrace / CLOUD (managed), or ingested by Darktrace / CLOUD from an existing configuration. Data is retrieved via API access granted during initial authentication.

The logs created are retrieved by Darktrace / CLOUD and translated into connection events in the Darktrace platform, where they are made available to core Darktrace analysis. Flow log monitoring provides a baseline of network event activity that is, by nature, fragmentary. Logs are pre-processed by the cloud provider and generated only for supported interactions, connection stages, or services. Known content limitations include, but are not limited to, lack of application protocol (AWS) and fragmentary DNS information (Azure).

Delays may be incurred where the cloud provider does not make flow log events available to Darktrace / CLOUD for processing and analysis in a timely manner. Delays of this nature, including any fixed latency, are the responsibility of the cloud provider.

1.2.4.3. Other Ingestion

Darktrace / CLOUD real-time detection can also receive and parse log data in syslog from external tools with additional configuration. Master, Unified View, physical Probe, and vSensor instances support syslog ingestion.

Darktrace offers a suite of Threat Intelligence and telemetry integrations where data may be retrieved from other compatible security tools or OSINT sources. The retrieval method for each integration is detailed in the corresponding integration documentation, which may be found on the Customer Portal. These alternative inputs (“Threat Intelligence Integrations” and “Telemetry Integrations”) are configured on the Darktrace Threat Visualizer “System Config” page.

Darktrace also provides a REST API for automation and a subset of relevant data input.

1.3. Darktrace Analysis

Darktrace analysis is multi-stage. The following outline highlights key elements of the process but is not intended to be an exhaustive list of all analysis performed during operation of the Darktrace platform, Threat Visualizer interface, or any underlying components.

Where source data is low quality, represents only a subset of all network activity, is delayed by external factors, or is otherwise incomplete, the quality of Darktrace analysis will be severely impacted. It is therefore imperative that data ingestion is configured comprehensively and to a high standard by the client. Analysis may also be impacted by overloading, such as that observed when traffic throughput significantly increases beyond scoped levels.

1.3.1. Asset Analysis

Darktrace / CLOUD amasses basic contextual information about assets from initial enumeration. Additional analysis is then performed to identify, derive, and associate further metadata to enrich the asset.

Darktrace / CLOUD will attempt to identify attached assets - other assets which are attached to, owned by, or possessed in some way by the asset under consideration – as well as associated assets. Analysis is also performed to identify asset access – assets that the chosen asset is permitted to access, or which are permitted to access the asset itself.

This analysis is performed using methodologies and techniques outlined under “1.2.2.2 Architecture Analysis” below.

Darktrace / CLOUD assesses the configuration of assets, policies, and the overall cloud environment against key guidelines drawn from various industry compliance frameworks and other assessment criteria. These additional criteria include general cloud-based best practices not unique to a specific framework and Darktrace-defined detections for insecure or unmaintained resource configuration.

Asset configuration (such as encryption, policies, or settings), access (network and permission-based), and status (e.g., runtime) are evaluated against pre-defined misconfiguration detections created by Darktrace analysts. Assets found to match known misconfigurations will be assigned a misconfiguration alert (see below).

1.3.2. Architecture Analysis

Architectures are data structures constructed by Darktrace / CLOUD from enumerated assets. Architectures are intended to represent interacting groups of entities in the cloud environment. The processes of discovery and analysis are entirely automated and are conducted at regular intervals to improve the accuracy of the representation of the active cloud environment.

The following multi-stage analysis utilizes a range of approaches including supervised machine learning, unsupervised machine learning, statistical analysis, and techniques derived from graph theory.

1.3.2.1. Architecture Construction Analysis

After Darktrace / CLOUD has enumerated all assets in the cloud environment, it will create connections between these assets using a first-stage analysis. This analysis simulates access restrictions for IAM (Identity and Access Management) and network policies, mapping the routes that exist between cloud resources, users, and roles. Other simulated connections or mappings may also be derived.

Once mapped, architecture analysis automatically assembles these assets and links into architectures - grouped collections of components that interact closely, representing connected systems and infrastructure in the cloud. Constructed architectures are then assigned auto-generated identifiers. Architectures are not static data structures; if an architecture significantly changes, for example due to a large number of new cloud resources, it may be remapped.

Architectures are not mutually exclusive structures - assets may be part of multiple architectures, or smaller architectures may be derived as part of larger, parent infrastructure.

1.3.2.2. Architecture Priority Analysis

After construction, each architecture is subject to prioritization analysis. Prioritization analysis considers factors including, but not limited to, the cost of the operation for the architecture, any model alerts or misconfigurations present on assets within the architecture, the type of assets present in the architecture (including service), the permissions assigned to entities in the architecture, any manual prioritization applied by an operator, or any other relevant metadata derived for assets within the architecture.

1.3.2.3. Permissions Analysis

In the course of enumeration and architecture construction, Darktrace / CLOUD performs permissions and access mapping analysis. This analysis results in pseudo-architecture structures representing all possible access between an identity and a cloud resource.

1.3.2.4. Attack Path Modeling

Darktrace / CLOUD utilizes a limited-feature version of Darktrace/Proactive Exposure Management attack path modeling to derive and display critical attack paths within architectures. The supported path types for analysis will differ between cloud providers and supported services.

1.3.3. Network Traffic Analysis

Analysis performed on network traffic prior to 'Pattern of Life' Analysis and Classification differs between ingestion modes.

1.3.3.1. Darktrace Deep Packet Inspection

Darktrace Deep Packet Inspection is performed on raw ingested network traffic by Darktrace virtual sensors (vSensors) to produce metadata for analysis by other platform components. The list of supported L4 - L7 protocols for inspection is regularly expanded, and analysis is continuously refined.

The output of Deep Packet Inspection is passed to the Darktrace model engine, to the core Darktrace "classifier" engine, and made available for display in the user interface and any other platform components such as the Cyber AI Analyst.

Deep Packet Inspection is performed only on network traffic ingested by Darktrace virtual sensors through a supported mode. Deep Packet Inspection is not performed on network data observed by flow logs.

1.3.3.2. Flow Logs

Flow logs are translated into standard format connection events for consumption by the Darktrace model engine, the core Darktrace “classifier” engine, and other subsequent components. Less in-depth analysis is performed on flow log data prior to submission to these components.

1.3.4. User Activity Event Inspection

Darktrace / CLOUD also retrieves identity data within the cloud environment information made available by third-party APIs (or other comparable methods) for analysis. Event and activity data is then parsed into a series of standard metadata fields and analyzed to identify the type of activity, any entities involved (e.g., users, files, or other resources), and to extract key information.

The output of this analysis is passed to the Darktrace model engine, to the core Darktrace “classifier” engine, and made available for display in the user interface and any other platform components such as the Cyber AI Analyst.

1.3.5. Darktrace ‘Pattern of Life’ Analysis and Classification

Network events, user activity events, connection data, and any other configured inputs are subject to Darktrace ‘pattern of life’ analysis. Darktrace will create individualized behavioral profiles for the network entities it observes and surface activity which is considered inconsistent with the expected norm. This baseline is derived from - but not limited to - an analysis of the behavior of the individual entity, analysis of one or more clusters created from similarly behaving entities, and many variable factors such as time or communication protocol. The ‘pattern of life’ data is continually updated in real-time and reflects the data that has been received, with a greater weighting to more recent data.

Analysis is performed in the “classifier” stage by a multitude of classifiers. This core analysis applies many approaches including unsupervised Machine Learning techniques such as Bayesian Meta-classification, techniques derived from graph theory and network analysis such as Node/Graph Centrality, approaches derived from statistical analysis such as Spectral Clustering, and many other techniques. The previous examples of utilized techniques provided are for illustrative purposes and should not be considered exhaustive.

Darktrace does not offer any capability to access underlying behavioral models or classification output. Darktrace provides the ability to utilize the output of this analysis using the *models* framework. Darktrace will create alerts to indicate anomalous activity which will be inserted into the “event log” within the Darktrace Threat Visualizer interface of the corresponding device, user, or entity. These unusual activity “notices” are also consumed by Darktrace real-time detection models.

1.3.6. Darktrace Real-Time Detection Models Framework

The Darktrace Real-Time Detection model engine is the logical framework within which events and the output of ‘pattern of life’ analysis is evaluated. Both network events and audit activity events described above submit data to the model engine for evaluation.

A model is used to define a set of conditions which, when met, will alert the system to the occurrence of a particular event or chain of anomalous behavior. The models framework leverages both the underlying

'pattern of life' detection and outputs from Darktrace Deep Packet Inspection, telemetry inputs, Darktrace / CLOUD, and other data sources. Output from the complex anomaly framework is available in accessible, building block format and can be combined with simple conditions and logical expressions to create tailored activity detection.

Custom models can be defined by operators of the Darktrace Threat Visualizer to meet any criteria; these are created and maintained in the Threat Visualizer "Model Editor" interface. Care should be taken when defining custom models to ensure that the number of alerts produced is not excessive and does not impact the system's ability to function, or otherwise degrade service. Similarly, modification of existing default models should not result in excessive activity or alter the logic to the extent that alerts criteria can no longer be met. The responsibility to ensure that models created or edited remain within these reasonable expected boundaries lies with the operator.

Default Darktrace models are focused on 'pattern of life' anomaly detection, potentially malicious behavior, and optional compliance issues. Darktrace will periodically update these standard supplied models - customers with Call-Home or Darktrace-hosted cloud instances will receive updates automatically; clients without automatic updates will receive all applicable model updates when Darktrace Threat Visualizer software is updated. The conditions for auto-update eligibility are described in the relevant product documentation.

Models may also be used to trigger actions within the Darktrace Threat Visualizer platform; the output of the Darktrace real-time detection model engine is described below.

1.3.6.1. Flow Logs & Real-Time Detection Models

Due to the fragmentary nature of network events derived from flow logs, not all Darktrace Real-Time Detection Models support detection solely from flow logs.

1.3.6.2. Autonomous Response Models

In default operation, Darktrace / CLOUD autonomous responses are triggered by model alerts from a specific subset of Darktrace models, categorized as autonomous response ("Antigena") models. Darktrace / CLOUD uses both standard Darktrace / NETWORK autonomous response models and bespoke models. Please refer to the Darktrace / NETWORK Autonomous Response specification for further information on these models.

Darktrace autonomous response models may directly look for specific behavior or for indicators identified by other models operating within the real-time detection environment. Darktrace autonomous response models typically fall into the second category - "meta-models" - which are triggered by an alert of another Darktrace Real-Time Detection model.

Any modification to an underlying real-time detection model which results in increased or reduced model alerts will subsequently impact autonomous response models and actions. Any under-activity or over-activity of the corresponding autonomous response model as a result of this type of modification is the responsibility of the operator.

Darktrace also provides the ability to limit actions created by these models to a minimum score threshold.

1.3.7. Cyber AI Analyst

Darktrace Cyber AI Analyst performs a meta-analysis upon the previous layers of analysis described. Please refer to Cyber AI Analyst in the Product Agnostic section below.

1.3.8. Cross-Capability / Cross-Coverage Area Analysis

Where other Darktrace ActiveAI Security Platform coverage areas are deployed, cross-platform analysis is performed by multiple components. Links are created for analysis purposes between network entities and entities modeled by other Darktrace real-time detection components such as, for example, user entities created by Darktrace / IDENTITY module monitoring known to be associated with a given network entity, or credential entries observed by Darktrace / NETWORK monitoring also observed as part of Darktrace / ENDPOINT.

Darktrace Cyber AI Analyst will also link together entities across different coverage areas in the creation of Cyber AI Analyst Incidents and Cyber AI Analyst Incidents Events. Darktrace Cyber AI Analyst may also retrieve additional contextual data during investigation from other components if deployed, such as the retrieval of associated emails from Darktrace/ EMAIL. Please refer to “Cyber AI Analyst” below.

The output from Darktrace components may be fed to Darktrace real-time detection as part of supplementary Threat Intelligence; Darktrace Attack Surface Management (ASM) can provide output of malicious asset identification into the “Watched Domains” list of Darktrace real-time detection to create model alerts (*formerly known as “model breach” alerts*) when a network entity accesses an ASM-identified domain.

The above provides illustrative examples of collaborative scenarios between platform components for reference, but is non-exhaustive.

1.4. Output

1.4.1. Processed Event Metadata

The metadata output of Darktrace Deep Packet Inspection, flow log analysis, and audited identity analysis is displayed in the Darktrace Threat Visualizer “Advanced Search” interface. Data retention for this output is on a rolling basis and is dependent upon a number of factors such as traffic makeup, and other operational components which contribute entries to Advanced Search. Most deployments can expect around 30 days retention. Darktrace Deep Packet Inspection metadata can be exported to external tooling by supported methods for longer retention.

1.4.2. Alerts

Darktrace / CLOUD real-time detection produces three primary alert types: Darktrace model alerts, misconfiguration alerts, and Darktrace Cyber AI Analyst Incidents.

Darktrace model alerts and misconfiguration alerts can be investigated within the Darktrace / CLOUD console, - a specialized user interface for Darktrace / CLOUD coverage - or exported to a compatible alert consumer.

Darktrace model alerts and Darktrace Cyber AI Analyst Incidents can also be investigated in the main Threat Visualizer interface, investigated in the Darktrace Mobile App, or exported to a compatible alert consumer.

Alerts are categorized for priority and filterable by multiple factors, allowing for customization of alerts displayed in Darktrace interfaces and those exported to external tools.

A secondary alert type - Darktrace System Status alerts - may also be generated where the Darktrace platform is experiencing degraded service due to health issues, invalid input, or where it is necessary to highlight changes to system administrators. Alerts include details of the originating host, the severity of

the event, and relevant links to investigate or resolve the issue. Notifications are sent when a system event becomes active and can optionally be sent on resolution.

Darktrace supports alert export in both industry-standard generic forms such as Syslog or Email and custom integrations with common tools such as Splunk and ServiceNow. Alert outputs (“Workflow Integrations”) are configured on the Darktrace Threat Visualizer “System Config” page. The details included in each external output may vary due to third-party restrictions on content length or supported formats.

1.4.2.1. Model Alerts

Darktrace model alerts are created as a result of Darktrace / CLOUD real-time detection models; when conditions for a model are met, a model alert can be created in addition to other possible model actions. Darktrace model alerts contain details of the conditions that were satisfied in order to trigger the alert, the entity which met those conditions (for example, a device, or another model alert), and a description with recommended action points.

1.4.2.2. Darktrace Cyber AI Analyst Incident Alerts

Darktrace Cyber AI Analyst Incident alerts are created when Darktrace Cyber AI Analyst identifies activity considered significant enough to highlight to operators. Darktrace Cyber AI Analyst Incidents may contain a single “Cyber AI Analyst Incident Event” or multiple linked findings aggregated together. External alerts are created when a new “Cyber AI Analyst Incident Event” is created, which may be associated with an existing “Cyber AI Analyst Incident”, or form an independent, new “Cyber AI Analyst Incident”. In the former case, mechanisms to identify the relationship between Cyber AI Analyst Incident Events are provided in the output.

In the Threat Visualizer and Darktrace Mobile App, the output of Darktrace Cyber AI Analyst analysis is aggregated on a per-Cyber AI Analyst Incident basis. New Cyber AI Analyst Incident Events added to existing Incidents will not produce independent alerts in these interfaces and are instead displayed as part of the existing Incident. Darktrace Cyber AI Analyst Incidents in these interfaces contain details of the activity identified, the Darktrace model alerts that triggered the initial Cyber AI Analyst investigation, the entities which performed the unusual activity (for example, devices), the investigation steps Cyber AI Analyst performed, why activity was linked together (if multiple Cyber AI Analyst Incident Events), and a human-readable summary of the finding.

1.4.2.3. Misconfiguration Alerts

Misconfiguration alerts identify assets which do not align with best practices or operational guidelines defined by cloud compliance frameworks. Each misconfiguration is created for an asset, and then many of the same misconfiguration are grouped together for ease of investigation.

Misconfigurations are categorized by priority and dynamically scored. Scoring may be impacted by factors including, but not limited to, the severity of the misconfiguration, the priority of the affected asset, and the overall priority of the architecture(s) the asset is a member of.

Misconfiguration alerts are intended to be resolved by user operators and will automatically resolve if the criteria are no longer found at the next enumeration of the cloud environment (e.g., due to correction, or the deletion of the associated asset). Options are also provided to mark misconfigurations as mitigated, accepted, and dismissed. Each state is defined in the relevant Darktrace / CLOUD documentation as found on the Customer Portal.

Once a user operator sets a misconfiguration alert's status to mitigated, accepted or dismissed via the user interface, or where a misconfiguration is created by Darktrace and deemed accepted, Darktrace / CLOUD will not generate further misconfiguration alerts for the relevant misconfiguration. Darktrace bears no liability for the misconfigurations marked mitigated, accepted or dismissed.

1.4.3. Autonomous Response

Darktrace / CLOUD operates two autonomous response approaches - "Darktrace / CLOUD core" and "Darktrace / CLOUD native." These approaches are not mutually exclusive; multiple approaches may be invoked as the result of a single model alert.

1.4.3.1. Darktrace / CLOUD Core Autonomous Response

Darktrace / CLOUD core creates automatic, network-level responses to anomalous connectivity detected by Darktrace / CLOUD real-time detection or when triggered manually by a Darktrace operator. Actions are applied by Darktrace vSensors or host-based osSensors.

Darktrace / CLOUD core operates under the same eligibility and action methodology as Darktrace / NETWORK autonomous response. Please refer to the description for this capability for further information on eligibility, application, and method.

Darktrace / CLOUD core does not support actions against connectivity monitored with flow logs. Darktrace / CLOUD core does not support "Integration Actions" taken through third-party network components.

1.4.3.2. Darktrace / CLOUD Native Autonomous Response

Darktrace / CLOUD native creates automatic responses using cloud providers' APIs; the range of actions available will therefore differ between cloud providers.

Darktrace / CLOUD native actions may be triggered by Darktrace / NETWORK autonomous response models, by bespoke Darktrace / CLOUD models, or manually by a Darktrace operator. At present, eligibility for Darktrace / CLOUD native actions is automatic.

Individual "inhibitors" - types of action taken - can be manually disabled from the Darktrace / CLOUD console. Where a disabled inhibitor is invoked, no action is taken - care should therefore be taken when modifying inhibitors to ensure Darktrace autonomous response actions are not prevented entirely.

1.5. Framework Assessment

Darktrace / CLOUD assesses the configuration of assets, policies, and the overall cloud environment against key guidelines drawn from various industry compliance frameworks and other assessment criteria. These additional criteria include general cloud-based best practices not unique to a specific framework and Darktrace-defined detections for insecure or unmaintained resource configuration. Scoring against these frameworks is surfaced in the Darktrace / CLOUD console.

Any assessment provided by the Darktrace / CLOUD console does not amount to a statement, warranty or attestation of compliance against any given compliance framework. Darktrace / CLOUD is not a compliance tool and Darktrace explicitly disclaims any and all liability for Customers' compliance under any applicable framework to fullest extent permitted by law.

1.6. Interface

The primary user interface of the Darktrace platform is the Threat Visualizer. The Darktrace / CLOUD Console is a purpose-built investigation interface for Darktrace / CLOUD that forms part of the Threat Visualizer.

Both the Darktrace / CLOUD Console and main Threat Visualizer interface include filterable alerts, detailed device information, and further investigation tools.

1.6.1. Threat Visualizer

The Threat Visualizer interface provides access to all major Darktrace user interfaces, consoles, and product views - it contains both investigation interfaces and administration interfaces. A list of the user interfaces that comprise the Threat Visualizer - and detailed information on how to operate these interfaces - is provided in the relevant technical documentation.

Alternative investigation views, access to review detailed metadata, report generation, and the ability to edit and create custom models (as described above) are provided within the main Threat Visualizer interface. Operators can also review the detected users and devices observed by Darktrace / CLOUD real-time detection (including any relevant contextual data), configure system settings, deploy integrations, review system health information, perform user management, and other administrative tasks.

1.6.2. Darktrace / CLOUD Console

The CLOUD Console includes interactive architecture visualizations, focused alert views for model alerts and misconfigurations, compliance overviews, and a queryable database of enriched asset data.

1.7. Usage metrics

Usage of Darktrace / CLOUD is calculated by reference to the two asset types of Identities and Workloads (as expanded on below), which, when taken together in the aggregate represent the Identity Workload count. Assets are weighted according to the computation required in analysis and the number of assets that make up a Workload or Identity will vary depending on the deployment.

Unless an asset is excluded by Darktrace, all asset types identified in the Customer environment will count toward the Identity Workload count. Darktrace / CLOUD is continuously expanding the range of supported cloud services and assets across all supported cloud providers. The most up to date list of asset types contributing to the Identity Workload count in a deployment can be made available on request.

Darktrace reserves the right to modify the asset types classified and modify any weighting assigned to individual asset types.

1.7.1.1. Identities

Assets enumerated as Identities by Darktrace / CLOUD include, but are not limited to: users with permissions to act within the applicable cloud provider environment, user-like asset types, and associated assets which impact IAM (Identity and Access Management), such as policies and entity management. Each user, user-like asset and associated asset will contribute toward the Identity count for usage metrics purposes.

1.7.1.2. Workloads

Workloads are assets in the cloud provider environment that can hold metadata, originate from a supported cloud provider service, and have not been explicitly excluded from enumeration (for example, default entities present in all environments with no relevant use).

Assets enumerated as Workloads by Darktrace / CLOUD include: virtual machines, compute functions, storage solutions, messaging streaming services, serverless instances, firewall rules, networking resources and other functional asset types.

All assets that are enumerated as assets are weighted according to their computational requirements. In expected operation, virtual machines and compute functions are more likely to amount to one Workload for usage metric purposes, whilst a varying number of serverless instances (and other functional asset types) will amount to one Workload.

1.8. Reporting

1.8.1. Auditing

The Threat Visualizer “Audit Log” records changes made by operators such as model alert acknowledgment; the audit log can be exported via Syslog for extended retention. User interactions with Darktrace / CLOUD autonomous response actions which alter the state (e.g., “activate,” “extend,” “clear”) are also recorded in the action history, accessible from the Darktrace Response Actions window. Users can also be compelled to provide a free-text justification when an action state is altered, which is also displayed where applicable.

1.8.2. Event Metadata

The output of Darktrace / CLOUD activity and network data retrieval and analysis is accessible in the Darktrace Threat Visualizer “Advanced Search” interface. Combined output from this processing, from Darktrace analysis, and from any actions performed by the platform automatically (such as tagging as a result of a model) are combined into logs which are displayed for each user.

Data retention for metadata output and combined log data is on a rolling basis and is dependent upon a number of factors such as hardware capability, event volume, and other operational components which contribute data to the platform. Most deployments can expect around 30 days retention of Darktrace / CLOUD metadata and general device/user activity data. Automatic removal of older event log data does not affect the storage or training of the machine learning ‘pattern of life’ data.

Darktrace also provides a REST API for automated retrieval of a subset of data output.

1.8.3. User and Device Entities

Cloud assets observed by Darktrace / CLOUD real-time detection to be recently active in the network environment are detailed on the Threat Visualizer “Device Admin” page. Subnets observed by Darktrace / CLOUD real-time detection in network connectivity are displayed on the Threat Visualizer “Subnet Admin” page. Metrics regarding data and event throughput are rendered on the Threat Visualizer “System Status” page.

1.8.4. Darktrace / CLOUD Console

All remaining metrics, alerts, interactions, and administration activities not present in the main Threat Visualizer workspace or dedicated administration pages can be found in the Darktrace / CLOUD Console.

2. Deployment Architecture

Darktrace / CLOUD deployment architecture comprises three main components: Darktrace ActiveAI Security Platform instance, hosted Darktrace / CLOUD infrastructure, and virtual sensors deployed in the cloud provider.

2.1. Darktrace (ActiveAI Security Platform) Instance

Darktrace / CLOUD can operate on Darktrace physical or cloud-hosted master instances, or on Unified View instances. Darktrace / CLOUD does not operate on physical probe instances, vSensors, osSensors, or other topology components not already referenced. Internet connectivity to the relevant endpoints - both those associated with the cloud provider and with relevant Darktrace services - is necessary for data retrieval.

2.2. Hosted Darktrace / CLOUD Infrastructure

Darktrace / CLOUD also comprises cloud-hosted components and infrastructure provisioned and maintained by Darktrace Operations in the Darktrace-owned public cloud environment.

All data generated from enumeration of cloud assets is isolated within a private namespace, stored in a region close to the customer in the Darktrace CLOUD environment. This data is encrypted at rest.

Darktrace / CLOUD utilizes industry-standard practices to segregate tenant data. All PII, processed and non-processed data is stored in isolated environments per client, with the exception of Graph databases which are fully multi-tenant. This multi-tenant architecture utilizes industry-standard practices for data segregation.

2.3. Virtual Sensors

Darktrace / CLOUD deployments also utilize standard Darktrace virtual sensors - vSensors, osSensors, and containerSensors - to extend traffic monitoring visibility across the cloud network environment.

Darktrace vSensors can process raw network traffic ingested directly from virtualized networking equipment, or mirrored to it by connected osSensor agents, and forward the resultant metadata onward to a connected Darktrace instance.

For more information about vSensors, osSensors and containerSensors please see the applicable passages of the Darktrace / NETWORK section.

3. Implementation

Darktrace / CLOUD is implemented through deployment of one or more physical or virtual components, including at least one Master instance. Darktrace / CLOUD hosted infrastructure is in AWS, and is available in different regions close to the Customer.

3.1. Darktrace Cloud for AWS

3.1.1. Deployment Process

In standard setup mode, Darktrace / CLOUD is deployed via commands entered into an AWS CloudShell by an AWS admin user. These commands use CloudFormation Stacks and StackSets to deploy the resources required into the AWS environment. The setup process can be performed for a single AWS account, or from the AWS Organization management account which will automatically include all child accounts. The latter is strongly recommended where possible.

Darktrace / CLOUD can also be deployed without CloudShell, by either deploying the CloudFormation Stack directly or using a Terraform file.

3.1.2. Deployment Scenarios

Darktrace / CLOUD can be deployed at the organization or account level.

In an organization-level deployment scenario, setup is performed by an authenticated admin user for the AWS Organization management account (root account). Where possible, resources are created at the organization account level to serve all child accounts.

In an account-level deployment scenario, setup is performed individually for each account by an authenticated admin user for that account. Steps must be repeated for each account intended for monitoring. Resources are created in each account and cannot be managed centrally.

3.1.3. Security Features

During installation, optional features are offered alongside the minimum required component (Darktrace / CLOUD Core). These optional features can be included during deployment or added at a later date.

3.1.3.1. CLOUD core (required)

The Darktrace / CLOUD Core component is required for Darktrace / CLOUD operation; this feature deploys the role necessary for resource enumeration and can also deploy (or upgrade) audit activity monitoring. At a minimum, this component will deploy an IAM role to the organization account or selected account.

Darktrace / CLOUD requires a CloudTrail logging to an accessible S3 bucket to monitor audit activity within the AWS environment. This monitoring can utilize an existing CloudTrail and S3 bucket, or in their absence it can create on CloudTrail and S3 bucket.

3.1.3.2. Network Flow Analysis (optional)

Network Flow Analysis enables the generation of VPC Flow Logs and Route 53 Resolver Query Logs, providing simple visibility into network activity for VPCs not covered by high-fidelity Darktrace Deep Packet Inspection. If included, the feature will deploy an IAM role to the selected account or the parent organization account (and child accounts) with permission to enable VPC flow log generation, to create S3 buckets to contain logs, and to delete resources created by the role. Existing flow logs bucket can be utilized as well.

The logs created are retrieved by Darktrace / CLOUD and translated into connection events in the Darktrace platform, where they are made available to core Darktrace analysis. 'Pattern of life analysis' can then be modelled for network assets observed in these flow logs, and events are accessible from both Advanced Search and the main Threat Visualizer workspace.

Devices modelled by this component will be placed into VPC-specific VLANs in the Darktrace Threat Visualizer and associated back to resources enumerated by Darktrace / CLOUD.

3.1.3.3. Costing (optional)

Darktrace / CLOUD can utilize AWS cost reporting to enhance the priority analysis of the resources and architectures it sees. This component enables the creation of Cost and Usage Reports by AWS, which are retrieved and processed by Darktrace.

Please note, costing components are always deployed in region us-east-1, regardless of the region selected during configuration or the primary region where assets are deployed. This requirement is enforced by AWS.

3.1.3.4. Autonomous Response (optional)

The Autonomous Response framework leverages the power of the 'pattern of life' developed across the platform to respond, contain, and neutralize emerging threats across the entire digital estate. The Autonomous Response component enables new, cloud-native Autonomous Response actions including role modification and security group restriction, ensuring targeted response for cloud assets.

3.2. Darktrace Cloud for Azure

3.2.1. Deployment Process

Deployment process is performed via Darktrace / CLOUD platform.

Darktrace / CLOUD is deployed via commands entered into an Azure CloudShell by a Global Admin with owner permissions over the tenant and management groups or subscriptions. The setup process can be performed on a subscription level, management group level or tenant level which will automatically include all subscriptions. The latter is strongly recommended where possible.

3.2.2. Deployment Scenarios

Darktrace Cloud for Azure can be deployed at the tenant or management group / subscription level. This is applicable to both the minimum requirement for Darktrace Cloud Security operation (Cloud Security Core) and to optional features such as Network Flow Analysis.

In a tenant-level deployment scenario, setup is performed by a Global Administrator with owner permissions over the tenant and tenant root management group. Components are then created within any child management groups or subscriptions under the tenant.

In a management group or subscription-level deployment scenario, setup is performed individually for each subscription or management group by a Global Administrator with owner permissions. Any new subscriptions that are created outside of the management group that you specified will not be included.

3.2.3. Security Features

3.2.3.1. CLOUD core (required)

The Darktrace / CLOUD Core component is required for Darktrace / CLOUD operation; this feature deploys the application registration necessary for resource enumeration. At a minimum, this component will deploy an IAM role to the management groups or selected subscriptions.

Darktrace / CLOUD requires a minimum P1 license activity, sign-in and audit logs retrieved by Graph API.

3.2.3.2. Network Flow Analysis (optional)

The Network flow analysis feature enables the generation of virtual network flow logs and optional tracking of DNS logging for agentless Network monitoring.

The Darktrace-Managed Flow Logs option will determine whether this module should automatically configure vNets for network flow analysis. This will involve creating Storage Accounts, Network Watchers, and VNet Flow Log resources in your subscription(s). Using this option is strongly recommended. If disabled, and Network Flow Analysis feature utilized, these resources must be manually created and maintained.

For more information, refer to the relevant section in the Customer Portal documentation.

3.2.3.3. Azure AD (Microsoft Entra ID) (optional)

Azure AD (Microsoft Entra ID) tracks user activity and sign-on metrics from the Graph API and the Microsoft Azure audit log endpoint.

3.2.3.4. 3.2.2.4 DNS Monitoring (optional)

The Enable DNS Logging option will determine whether the module should try and pull DNS logs from Azure Sentinel. Please note that you will need a Windows DNS server in order to use this setting.

For more information, refer to the relevant section in the Customer Portal documentation.

3.2.3.5. Autonomous Response (optional)

Autonomous Response enables mitigations of suspicious activity within your environment by creating an app registration with higher privileges.

3.3. Data checks for Darktrace / CLOUD

To confirm successful installation and the CLOUD platform is accessible, post installation, data checks should be carried out. Details of the necessary checks for each instance are provided in the Customer Portal documentation.

3.4. Darktrace virtual sensors

In scenarios where L4-L7 traffic inspection is required, network traffic can be sent via a Darktrace vSensor in one of two encrypted communication modes. vSensors can be deployed in a traffic-mirroring scenario, or with up to 255 osSensor agents (per vSensor). Darktrace osSensors can be installed on devices running Windows, supported Linux distributions and any Linux environment running the Docker engine. ContainerSensors can be deployed on Linux and Windows Kubernetes hosts to integrate with containerized environments.

Darktrace vSensors will automatically update if granted access to the Darktrace update packages infrastructure (packages.darktrace.com or packages-cdn.darktrace.com). The update schedule for the Darktrace osSensor offered in container format is defined by the client during configuration. Other Darktrace osSensor agent formats do not automatically update and must be managed by the client.

3.4.1. Deployment process for Darktrace vSensors

The vSensors need to be configured by the customer with a single admin interface. The admin interface is utilized for management and ingestion of traffic.

The vSensor supports VXLAN, GENEVE, ERSPAN traffics type I and type II, as well as GRE. It also supports ingestion of traffic from native mirroring technologies, such as AWS VPC Traffic Mirroring and GCP Packet Mirroring, or multiple connected osSensors. The osSensor agent is installed on each customer device where visibility is desired and monitors all the network traffic to/from configured interfaces of that device; the monitored traffic is then sent to the vSensor for analysis. osSensors utilize host resources to forward traffic, so should only be installed where it is not possible to retrieve traffic through other means.

Deployment documentation is provided on the Darktrace Customer Portal for each product.

3.4.2. vSensor Deployment Methods

Darktrace vSensors for AWS can be deployed as standalone image, QuickStart using CloudFormation or via a Terraform file. The deployment guides for each method are available in the Customer Portal.

3.4.3. vSensor Health Check

Health checks must be performed by the customer to confirm that the vSensor is running and has connectivity with the associated master and osSensors (if applicable). Key areas to test to verify health include: (i) that the vSensor virtual machine is running; (ii) verification of incoming packets; and checking for vSensor overload.

3.4.4. Deployment process for Darktrace osSensors

The Darktrace vSensor coordinates with the osSensors associated with it, ensuring traffic is captured only once when osSensor devices communicate to each other. Each osSensor registers with a vSensor using a shared HMAC token which should be supplied to both ends. It is recommended that the associated vSensor is configured in advance of osSensor setup, in order to ensure the necessary HMAC token and IP Address for the vSensor have been collected.

The osSensor is deployed as a package installed on the customer server. Each osSensor registers with a vSensor using a shared HMAC token which should be supplied to both ends. The associated vSensor should be configured in advance of osSensor setup, in order to ensure the necessary HMAC token and IP Address for the vSensor have been assigned.

The Darktrace vSensor coordinates with the osSensors associated with it, ensuring traffic is captured only once when osSensor devices communicate to each other. The installation guides for each available formats of osSensor can be found on the Customer Portal.

3.4.5. osSensor Deployment Methods

Darktrace osSensors for AWS can be deployed manually or using a package manager tools (e.g. SSM, Ansible, Jenkins and others). The deployment guides for each method are available in the Customer Portal.

3.4.6. osSensor Health Check

It is recommended that the Customer checks: (i) that the osSensor service is running and that there is connectivity with the associated vSensor; (ii) how many osSensors are running; and (iii) for vSensor overloading.

Further information on the testing commands can be found in the Customer Portal. It is a Customer responsibility to run periodical health checks to ensure the osSensor service is still running as optimally as initially setup.

3.4.7. Deployment process for Darktrace containerSensors

The containerSensor requires two environment variables set: the IP or Hostname of the connected vSensor and a HMAC token to authenticate communications to the vSensor. The HMAC token used is the same as that for osSensor to vSensor communication. More information on the deployment process can be found in the Customer Portal.

3.4.8. containerSensor Deployment Methods

Darktrace containerSensors for AWS can be deployed via kubectl or helm package manager. The deployment guides for each method are available in the Customer Portal.

3.4.9. containerSensor Health Check

It is recommended that the Customer checks: (i) that the containerSensor service is running and that there is connectivity with the associated vSensor and (ii) how many containerSensors are running.

Further information on the testing commands can be found in Customer Portal. It is a customer responsibility to run periodical health checks to ensure the containerSensor service is still running as optimally as initially setup.

4. Administration

General administration of the Darktrace deployment is performed in the Threat Visualizer interface:

- The majority of configuration, including administration of system settings such as proxies, authentication configuration such as LDAP and SAML SSO, deployment of alert and threat intelligence integrations, and other administrative tasks, are performed on the “System Config” page.
- User and group management, including the assignment of data visibility and permissions, is performed on the “Permissions Admin” page.
- Device administration, such as the setting of labels, application of tags, and altering of priority, is performed on the “Device Admin” page.
- Management of subnets observed by the Darktrace / CLOUD real-time detection instance is located on the “Subnet Admin” page.
- System health information and system alert resolution is performed on the “System Status” page.

For physical Darktrace appliances of any role, additional administrative tasks such as interface configuration, setting of host variables, and Call-Home configuration may be performed in the appliance console, accessible over SSH. For cloud-hosted instances, these administrative actions are managed by Darktrace operations directly and no access to the console is available.

Darktrace physical and cloud instances are each seeded with random passwords and two-factor authentication secrets at build time. These initial secrets are stored by Darktrace. Darktrace will provide the client with username/password combinations granting access to the Darktrace Threat Visualizer interface and, for physical instances, the SSH administration console (“Darktrace Console”); these passwords can be optionally changed by the client at any time.

Access to the underlying backend systems of the Darktrace instance is reserved for Darktrace only.

Darktrace vSensors only provide access for administration tasks via the management console. osSensor, TSA, and DAA agents do not provide administration interfaces and are instead managed with configuration files and/or dialogs during installation.

Detailed information about individual administrative tasks is outlined in the relevant documentation for the deployment of the specified component or administrative action intended to be performed.

4.1. CLOUD Console Administration

A subset of additional configuration options specific to Darktrace / CLOUD operations are offered from a dedicated status and administration menu within the Darktrace / CLOUD console. Documentation for the installation, operation and administration of Darktrace / CLOUD console is provided on the Darktrace Customer Portal (<https://customerportal.darktrace.com/product-guides/main/cloud-security-ui-universal-items>).

4.2. 4.2 Darktrace / CLOUD Quick Setup Wizard

Darktrace provides a “Quick Setup Process” (hereafter QSP) for Darktrace / CLOUD. This process guides users through enabling Darktrace / CLOUD for either AWS or Azure . This process offers both templated operating scenarios and granular configuration steps for each deployment aspect.

The QSP is intended for standard Darktrace / CLOUD deployment and operation. The QSP does not automatically deploy resources for Darktrace / CLOUD in AWS or Azure - these elements must be rolled-out by selecting Run Script in CloudShell, clicking Deploy with CloudShell Script, and then copying and pasting the Darktrace provided command into an authenticated AWS or Azure CloudShell. This command will download and run a shell script within the authenticated account and verify that the resources required for operation are created in your AWS or Azure environment. This is the recommended setup.

Darktrace/CLOUD for AWS can also be deployed without CloudShell, by either deploying the Darktrace provided CloudFormation Stack directly or using a Terraform file. These are more advanced setups that are described in full in [Deploying Darktrace/CLOUD](#).

4.3. 4.3. Administration of Darktrace / CLOUD Modules

The Darktrace /CLOUD QSP will automatically configure the relevant Darktrace / CLOUD Modules (AWS, AWS Network Flow Analysis, Azure, Azure Network Flow Analysis) on the Darktrace “System Config” page based on the input configured during the QSP. Each module has an individual entry with module-specific configuration settings.

Users can modify monitored accounts or subscriptions, flow log inclusion and exclusions, and authenticate or reauthenticate - the module from this location. It is a Customer responsibility to ensure that each module is configured correctly to ensure correct operation and interactivity with the associated third-party platform. These are more advanced setups that are described in full in [Ways to Deploy the Darktrace/CLOUD AWS Module](#) and [Deploying the Darktrace/CLOUD Azure Module](#).

4.4. De-Commissioning

The process for removing a Darktrace / CLOUD module will differ between third-party platforms.

Modules can be “de-authenticated” from the Darktrace System Config page - this will remove any authentication information used by the Darktrace ActiveAI Security Platform to contact the third-party platform. After de-authentication, any components created during the authentication process such as OAuth app registrations, service principals, API keys, or other authentication methods can be removed or deleted by the client from the relevant third-party platform. Darktrace strongly recommends that the client remove any remaining components when Darktrace / CLOUD module service is ended.

5. 5. Requirements

For organizations who already possess a Darktrace deployment such as Darktrace/NETWORK or Darktrace/ENDPOINT, this Darktrace master instance must be permitted to contact the Darktrace / CLOUD environment over 443/SSL. Your Darktrace / CLOUD environment will be created in an appropriate cloud region; details of the hostnames associated with Darktrace / CLOUD infrastructure can be found in Darktrace / CLOUD Region Availability and Security. The master instance also requires DNS to resolve these domains. For Darktrace cloud-hosted masters, Darktrace operations will ensure this access is in place.

If your organization is subject to restrictions on where enumeration data should be held (such as the GDPR Legislation in the European Union), please inform your Darktrace representative or note this restriction when submitting information on the Customer Portal.

5.1. Darktrace / CLOUD Region Availability and Security

If the customer organization is subject to regional or geographic restrictions on data flow that would prevent use any of the following cloud regions, the customer must disclose this to Darktrace.

All data generated from enumeration of cloud assets is isolated within a private name space and stored in an AWS region close to the customer in the Darktrace/CLOUD environment. This data is encrypted at rest. Access to the Darktrace/CLOUD interface is governed by user permissions. This is managed in the Permissions Admin page of the Darktrace Threat Visualizer interface.

5.1.1. AWS Hosting

Darktrace can offer AWS cloud-based deployments hosted in Europe (AWS region “eu-west-1” or “eu-west-2”), the United States (AWS regions “us-west-1” and “us-west-2”), Canada (AWS region “ca-central-1”), Singapore (AWS region “ap-southeast-1”) or Australia (AWS region “ap-southeast-2”).

5.1.2. Azure Hosting

Darktrace can offer Azure cloud-based deployments hosted in Europe (Azure region “UKSouth” or “WestEurope”), the United States (Azure region “EastUS”), Canada (Azure region “CanadaCentral”), South East Asia (Azure region “SoutheastAsia”), Japan (Azure region “japaneast”) or Australia (Azure region “AustraliaEast”).

5.2. Requirements of Darktrace / CLOUD for AWS

5.2.1. Pre-Setup Requirements for Darktrace / CLOUD for AWS

Darktrace Cloud Security features can be deployed at the organization or account level. This is applicable to both the minimum requirement for Darktrace Cloud Security operation (Cloud Security Core) and to optional features.

5.2.1.1. Organization-Level Authentication (recommended)

If you are authenticating at the organization level, sign into the management account for the parent AWS organization (that contains all organizations or subsidiary accounts) as an admin user. This user must have permission to create CloudTrails, Cloudformation Stacks and StackSets in the management account and child accounts.

Please refer to the following AWS documentation if you are not familiar with AWS Organizations delegated administrator permissions:

- [“AWS CloudFormation StackSets and AWS Organizations”](#)
- [“Organization delegated administrator \(CloudTrail\)”](#)

For organization-level deployment, also ensure Trusted Access is enabled to CloudTrail and StackSets in the AWS Organizations console. If you are not familiar with AWS Organization Trusted Access, please refer to the following AWS documentation:

- [“Enabling trusted access with CloudTrail”](#) in [“AWS CloudTrail and AWS Organizations”](#)
- [“Activate trusted access with AWS Organizations”](#)

5.2.1.2. Account-Level Authentication

If you are authenticating a single account, sign into this account as an administrative user with permission to create CloudTrails, and to deploy Stacks.

5.2.2. Selecting Cloud Security Features

During installation, optional features are offered alongside the minimum required component (Cloud Security Core). These optional features can be included during deployment or added at a later date. Please see Adding Additional Features to an Existing Darktrace Cloud Security Deployment for further details on how to add optional settings after configuration.

5.2.2.1. Cloud Security Core (required)

The Cloud Security Core component is required for Darktrace Cloud Security operation; this feature deploys the role necessary for resource enumeration and can also deploy (or upgrade) audit activity monitoring. At a minimum, this component will deploy an IAM role to the organization account or selected account.

Darktrace Cloud Security requires a CloudTrail logging to an accessible S3 bucket to monitor audit activity within the AWS environment. This monitoring can utilize an existing CloudTrail and S3 bucket - such as that created for the Darktrace DETECT/Cloud AWS module at previous date - or deploy these monitoring components afresh.

It is not recommended to deploy a new CloudTrail if logging is already enabled, as costs are incurred when multiple CloudTrails log the same management activity. Please refer to the AWS resource on Trail configuration for more information.

5.2.2.2. Network Flow Analysis (optional)

Network Flow Analysis enables the generation of VPC Flow Logs and Route 53 Resolver Query Logs, providing simple visibility into network activity for VPCs not covered by high-fidelity Darktrace Deep Packet Inspection. If included, the feature will deploy an IAM role to the selected account or the parent organization account (and child accounts) with permission to enable VPC flow log generation, to create S3 buckets to contain logs, and to delete resources created by the role.

During the setup process, the customer will be required to log into AWS as an Admin User or Delegated Admin User in the Organization Management Account or desired Individual accounts to authenticate Darktrace / CLOUD with the AWS environment. Darktrace uses External ID based Role Assumption. This mechanism is recommended by AWS for granting third party access. This process entails creating a role where the principal is a 3rd party AWS account, with the requirements to assume this role being that you must pass a known, and correct External ID.

5.3. Requirements of Darktrace / CLOUD for Azure

5.3.1. Pre-Setup Requirements for Darktrace / CLOUD for Azure

The following must be performed or retrieved before proceeding with the deployment process.

Darktrace Cloud Security features can be deployed at the tenant or management group / subscription level. This is applicable to both the minimum requirement for Darktrace Cloud Security operation (Cloud Security Core) and to optional features.

5.3.1.1. Global (Tenant-wide)

In a tenant-level deployment scenario, setup is performed by a Global Administrator with owner permissions over the tenant and tenant root management group. Components are then created within any child management groups or subscriptions under the tenant.

If selected, tenant properties must allow the tenant 'global access'.

5.3.1.2. Tenant Flow Requirements

5.3.1.2.1. Pre-Authentication as Global Administrator and Owner

The Tenant setup process must be performed by an Azure Global Administrator user with owner permissions over the tenant and tenant root management group. More information on this can be found in ["Elevate access to manage all Azure subscriptions and management groups"](#). Darktrace strongly recommends signing into the Azure portal with this user prior to steps outlined in ["Deploy Darktrace/CLOUD for Azure"](#). This will ensure that the Azure Cloud Shell can be launched successfully.

The user who performs the deployment process must also be permitted to manage all child subscriptions and management groups. How to confirm this access is in place is outlined below under "Grant Permission to Manage Child Resources".

5.3.1.2.2. Grant Permission to Manage Child Resources

To ensure that Darktrace/CLOUD monitoring can be deployed to all management groups and subscriptions, the authorizing user must be permitted to manage these resources from the tenant level.

How to provide this access is outlined under "Perform steps at root scope" in the Microsoft guide ["Elevate access to manage all Azure subscriptions and management groups"](#).

This access must be provided prior to deployment.

5.3.1.2.3. Tenant Primary Domain

Both deployment flows require your tenant domain. It can be retrieved from Azure by following the steps under "Find the Microsoft Entra tenant ID and primary domain name" in the Microsoft guide ["Locate important IDs for a user"](#).

5.3.1.3. Individual (Subscription or Management Group-level)

In a management group or subscription-level deployment scenario, setup is performed individually for each subscription or management group by a Global Administrator with owner permissions. Any new subscriptions that are created outside of the management group that you specified will not be included. To setup any management groups or subscriptions that were not covered during the initial roll-out, the full setup must be repeated. Resources are created in each individual subscription or management group and cannot be managed centrally.

5.3.1.3.1. *Individual Subscriptions or Management Groups Flow Requirements*

The Individual Subscriptions or Management Groups setup process must be performed by an Azure Global Administrator user with owner permissions over the subscriptions and/or management groups intended for monitoring.

Darktrace strongly recommends signing into the Azure portal with this user prior to steps outlined in "[Deploy Darktrace/CLOUD for Azure](#)". This will ensure that the Azure Cloud Shell can be launched successfully.

For Individual Subscriptions or Management Groups setup, the ID of the chosen subscription or management group will also be required during configuration.

Both deployment flows require your tenant domain. It can be retrieved from Azure by following the steps under "Find the Microsoft Entra tenant ID and primary domain name" in the Microsoft guide "[Locate important IDs for a user](#)".

5.3.2. Selecting Cloud Security Features

During the setup process, Customer will be required to log into Azure Portal as a Global Administrator to grant permission for Darktrace to monitor the Azure environment. Darktrace / CLOUD for Azure requires two applications: "Darktrace/CLOUD" (core operation) and "Darktrace Flow Analysis" (network traffic). The "Darktrace / CLOUD Azure" and "Darktrace Autonomous Response" modules are deployed separately. Roles are then assigned to the service principals created by these applications.

5.3.2.1. *DNS Monitoring (optional)*

If your virtual networks make use of a Windows DNS Server, Darktrace can extract DNS information from it. This will allow hostnames to be added to the network logs. If DNS is not configured, the Azure Network Flow Analysis module will still function, however hostnames will not be available.

Microsoft Sentinel must be installed and enabled. A Sentinel Workspace must also be in place.

The virtual network must be configured to use a Windows DNS server. This can be either through a Windows DNS server installed on an Azure Virtual Machine or through an on-premises Windows DNS server with the [Azure Arc](#) agent installed.

Access to 'api.loganalytics.azure.com' is required for the DNS component of Azure flowlogs to function. Please ensure this is permitted by any intermediary firewalls.

Please refer to [Configure Azure DNS Logs for Darktrace/CLOUD Enrichment](#) for the required steps.

Any further "Prerequisites" are outlined in [Stream and filter data from Windows DNS servers with the AMA connector](#)".

6. Customer Responsibilities

Notwithstanding any additional responsibilities set out in any other section of this Product Specification, and without limitation, Customer's responsibilities are to ensure that:

- Darktrace / CLOUD has, and maintains, visibility over all cloud accounts and cloud-hosted networks to be monitored.

- Where possible, Darktrace / CLOUD is deployed at the root level of the relevant cloud provider – root management group/tenant (Azure) or organization (AWS) - to ensure the fullest possible visibility.
- If not deployed at the root level, that Darktrace / CLOUD is deployed into all accounts at initial authorization and that authorization is repeated for any new cloud accounts created within the provider after this date. Darktrace / CLOUD cannot enumerate or monitor accounts outside its authorized scope.
- No configuration is in place within the cloud environment to partially or fully prevent Darktrace / CLOUD from operating. Examples include the presence of restrictive AWS Service Control Policies (SCPs) or Azure Deny Assignments.
- Darktrace / CLOUD has been configured and authorized according to Darktrace instructions. This includes ensuring that all permissions requested by Darktrace / CLOUD are granted, whether during an authentication flow such as OAuth, granted separately as permissions or roles in the cloud provider, or granted/provided in another comparable deployment process.
- If advanced deployment flows are used (for example, Terraform), all necessary components have been configured and authorized according to Darktrace instructions.
- Authentication is performed by - or granted via - a user with the appropriate permissions as outlined above.
- Recommendations or requirements placed on authorizing users are due to technical necessities. Darktrace cannot guarantee that OAuth grants, API keys, or other authentication methods performed by users without the appropriate permissions will result in functional operation.
- Darktrace / CLOUD retains authentication across all cloud accounts and, if lost, is re-authenticated in a timely manner. Unauthorized deployments are unable to continue with monitoring or Darktrace autonomous response.
- Modifications are not made to the cloud provider configuration or licensing that would prevent operation. Examples include the removal of auditing within the platform (such as via AWS CloudTrail), removal of required events from scope (for example, in Azure Event Hubs), removal of licenses required to contact specific APIs, removal of permissions from associated API keys, or any other configuration alteration that would interrupt service operation.
- If adding additional cloud accounts or migrating authentication mode, Darktrace / CLOUD is re-authenticated and permissions are granted in line with Darktrace recommendations. If necessary, remove any prior configuration before proceeding to change deployment mode.
- Darktrace / CLOUD autonomous response is enabled globally, in the deployment settings, and in any other applicable locations to ensure actions can be taken. Deployment configuration must also not preclude the taking of reasonable actions, for example, where all users are placed in an “immune” state or all inhibitors disabled.
- Darktrace / CLOUD is granted necessary autonomy to enact response actions, or, where human confirmation mode is partially or fully in place, ensure that actions are activated by human operators in a timely manner.
- Darktrace / CLOUD autonomous response has - and maintains - reachability to all devices intended for autonomous response actions. Where reachability is impeded, it falls within the customer’s responsibility to ensure that reachability is restored.
- Devices or user entities are opted appropriately into Darktrace / CLOUD autonomous response using tags.

- Alterations are not made to Darktrace / CLOUD autonomous response models or models utilized by Darktrace / CLOUD autonomous response that would significantly impede the ability to create targeted actions.
- Darktrace / CLOUD autonomous response actions are not created excessively by custom models or existing models modified to contain actions.
- Network traffic provided to Darktrace is of suitable quality and does not contain duplication, fragmentary data, or is incomplete in any fashion (for example, unidirectional). Where flow logs are utilized, ensure that flow logs contain all possible data made available by the cloud provider, including performing any additional required configuration.
- End-user access is managed appropriately and within recommended best practice.
- System health issues are monitored using the provided tools, and if alerts arise, addressed and rectified in a timely manner.

Role	Activity/Responsibility	
R - Responsible	The party is responsible for implementation of the activity; owns the problem or project	
A - Accountable	Right to make decisions, signs, and improves work	
C – Consulted	Has information, resources, and/or capacities necessary to assist the job	
I – Informed	Must be informed of the results, but does not need to be consulted	
Activity/Responsibility	Darktrace	Customer
Implementation, Administration & Maintenance		
Product Implementation, Administration and Maintenance	CI	RA
Provision of guides to support Implementation, Administration & Maintenance	RA	CI
Product Licensing	RA	CI
Implement Darktrace / CLOUD at a level that will provide highest visibility	CI	RA
Maintain coverage, including visibility, tagging and eligibility within product license	CI	RA
Monitor and resolve system status alerting	CI	RA
Performing regular CLOUD Response testing to ensure continued	CI	RA

reachability		
Ongoing maintenance of Autonomous Response configurations	CI	RA
Maintain & modify traffic feeds as required	CI	RA
Make software components available as detailed in customer contract terms	RA	CI
Download software from Darktrace infrastructure	CI	RA
Access and deploy new virtual machines to support sensor components	CI	RA
Modify firewall rules to enable communication between Darktrace master and sensor components	CI	RA
Modify configuration within the cloud environment that could otherwise partially or fully prevent Darktrace / CLOUD from operating	CI	RA

7. Considerations

Due to the variation in cloud providers, considerations will differ between individual providers. Please refer to the deployment material for the specific platform for further information.

Flow log monitoring provides a baseline of network event activity but is limited in content and, in most cloud providers, incurs a minimum latency. Known content limitations include, but are not limited to, lack of application protocol (AWS) and fragmentary DNS information (Azure). Darktrace / CLOUD core autonomous response is also not compatible with flow log monitoring; full traffic information is required.

If not deployed at the root level, Darktrace / CLOUD is limited to the accounts authorized by the client during initial authentication. Those which are not included in coverage cannot be enumerated, monitored, or included in autonomous response. Any new cloud accounts created within the provider after initial authentication must also be added by the client operator.

Limitations may exist where deployment at the root level (or equivalent) does not provide Darktrace / CLOUD with the required permissions to view all accounts due to the configuration of the cloud provider (for example, where no management groups are in use in Azure). Responsibility to ensure coverage remains comprehensive in this scenario lies with the client operator.

Darktrace / CLOUD analysis and monitoring is delivered by interaction with cloud provider APIs and supported integration methods. These APIs, services, and methods may incur costs within the cloud provider. Mitigations to reduce these costs are recommended where possible (for example, the reuse of existing CloudTrail logs), but are not exhaustive. Any costs incurred in the cloud environment as a result of Darktrace / CLOUD operation remain between the customer and the cloud provider.

If latency exists between the Darktrace ActiveAI Security Platform instance and Darktrace / CLOUD hosted infrastructure (for example, due to limited outbound bandwidth), or between the Darktrace ActiveAI Security Platform instance and the cloud provider (for example, due to delays in the cloud provider processing of requests), Darktrace / CLOUD Autonomous Response actions may incur delays.

The responsibility to ensure that the Darktrace ActiveAI Security Platform instance has an appropriate level of outbound access and bandwidth lies with the client operator.

The responsibility for any delays incurred between instruction by Darktrace / CLOUD and action occurrence in the cloud provider lies with the cloud provider.

Delays may be incurred where the external platform does not make events available to Darktrace / CLOUD for processing and analysis within the expected time frame. Delays of this nature are the responsibility of the cloud provider. Latency between event occurrence and when it was made available to Darktrace / CLOUD is indicated in the event metadata within the Threat Visualizer.

Darktrace / CLOUD autonomous response may be limited, or prevented from action entirely, if all relevant inhibitors have been manually disabled by the client operator.

Where delays are incurred receiving events from the cloud provider, Darktrace / CLOUD may be unable to take Autonomous Response actions in a timely manner due to the latency between event creation and its ingestion and processing by Darktrace / CLOUD.

Individual vendors may place rate-limiting restrictions on the APIs and methods utilized by Darktrace / CLOUD. Please refer to the individual provider documentation for further information.

Darktrace / CLOUD monitors and retrieves data from cloud providers using the APIs and integration methods presently offered by those platforms. If this coverage or available capabilities are modified or revoked by the third-party, Darktrace will endeavor to maintain continuity but is ultimately limited in capacity to restore service.

Darktrace Autonomous Response is subject to the restrictions placed by cloud providers on API capabilities. This can result in specific user types being ineligible for actions, specific circumstances impacting the efficacy of actions, and the requirement for additional configuration (in both Darktrace and/or the cloud provider).

A known consideration impacts Darktrace / CLOUD native actions against ongoing connections; due to limitations in API capabilities, ongoing connections cannot be ended. Only newly initiated connections can be prevented. This consideration does not impact Darktrace / CLOUD core actions taken by virtual sensors (osSensor or vSensor). Any further considerations are outlined in the individual provider documentation, where possible.

TCP RST packets created by Darktrace / CLOUD core autonomous response are typically unable to cross stateful boundaries and may also be limited in traversal by network routing components. Darktrace provides best practice guidance to mitigate these limitations.

The operator must ensure that guidelines are followed to provide routing to all locations. This includes permitting the traversal of spoofed packets between cloud environments or, where necessary, the deployment of additional Darktrace virtual sensors to ensure full coverage.

Changes made to cloud provider capabilities or APIs are outside the scope of Darktrace control and may be introduced at any time. Changes to the known considerations due to alterations by vendors of these platforms may arise at any time.

Darktrace / CLOUD may be reliant upon specific vendor licenses to retrieve the events necessary for operation or to take autonomous responses. If the required licenses are not present or are revoked, degradation or interruption of service will occur.

Where source data is low quality, represents only a subset of all cloud activity (network and/or activity), is delayed by external factors, or is otherwise incomplete, the quality of Darktrace analysis will be severely impacted. It is therefore imperative that data ingestion is configured comprehensively and to a high standard by the client. Analysis may also be impacted by overloading, such as that observed when traffic throughput significantly increases beyond scoped levels.

Where a Darktrace environment undergoes significant system load - typically due to exceeding recommended operational limits - "High performance mode" may be activated. This mode disables a subset of high-load 'pattern of life' classifiers to reduce CPU/memory usage within acceptable limits.

Environments which have not been adequately configured to track devices by the most appropriate, consistent factor, or where tracking has not been updated in line with network changes, will be unable to develop long-term 'pattern of life' behavioral profiles for unusual activity detection.

In the event of decommissioning, Darktrace / CLOUD-created resources should be removed from the cloud provider by the client operator at the point of decommissioning. This includes the removal of any Service Principals or application registrations, removal of IAM roles and policies, decommissioning of any functions or storage, and disabling of any logging rules.

8. Roles

Customer Role	Responsibility
Darktrace Product Owner / Project Manager	Coordinates Customer resources as necessary. Serves as the primary point of contact between Customer and Darktrace. Drives communication from the Customer side. Serves as the point of escalation for issue resolution and service-related issues.
Users of Darktrace Platform	Leverages Product functionality; monitors & actions output. Provides feedback to other Customer & Darktrace roles
Cloud/System Administrator	Ensures Darktrace/CLOUD coverage is maintained. Performs regular checks of data inputs. Responds accordingly to System Status alerts to ensure optimal traffic quality and delivery. Maintains call-home connectivity with Darktrace (where required).
Customer Portal Primary User	Manages customer portal access and contact information for all other customer roles. Ensures that Service Contacts are verified.

Darktrace Role	Responsibility
Customer Success Manager	Oversee customer's experience using Darktrace. Act as customers' sponsor & escalation path internally within Darktrace.

Darktrace Customer Support	Provides Support Services as per Master Services Agreement
Account Executive	Facilitate commercial matters between Customer and Darktrace
Solutions Engineer	Provides technical expertise to support Account Executive and Customer Success Manager, regarding expansions or changes of coverage.
Professional Services Engineer	Provides professional services to customer if subscribed, typically implementation support at start of contract

Darktrace / Forensic Acquisition & Investigation

1.	Introduction	132
2.	Product Features	133
2.1.	High Level Summary	133
2.2.	Data retrieval and ingestion	133
2.3.	Analysis.....	137
2.4.	Output	138
3.	Deployment Architecture	139
3.1.	Customer hosted	139
3.2.	SaaS hosted	141
4.	Implementation	142
4.1.	Customer hosted	142
4.2.	SaaS hosted	142
5.	Administration.....	142
5.1.	Settings	142
5.2.	Teams	143
6.	Customer Responsibilities	144
7.	Considerations.....	144
8.	Assumptions	145

1. Introduction

Darktrace / Forensic Acquisition & Investigation is a market leading cyber security solution for forensic data capture and analysis of both cloud-based and on-premise resources. This document is to be read in conjunction with the Darktrace Master Service Agreement which governs the usage of the Darktrace Product Offering. References to “Customer” throughout this document should be read to refer to the entity that is the owner of the software subscription and is ultimately responsible for its operation, whether as end-user or service provider.

Darktrace / Forensic Acquisition & Investigation was previously named CADO and sold by Cado Security Limited. Cado Security Limited have been acquired by Darktrace and this specification applies for Customers who have purchased Darktrace / Forensic Acquisition & Investigation under any other naming convention and regardless of the vendor or reseller.

2. Product Features

2.1. High Level Summary

Darktrace / Forensic Acquisition & Investigation (“Darktrace / FAI”) automates the forensic data capture, processing, and analysis of cloud-based resources (compute, containers, serverless, object storage, and logs), as well as on-premise systems – the latter through integrations with cloud native detection providers and XDR platforms (from hereout referred to collectively as “detection providers”). Darktrace / FAI uses a variety of data acquisition methods, however the main types are full volume capture, as well as a more lightweight triage capture tool.

There are two deployment options for the Darktrace / FAI product:

(1) Customer hosted – the product is deployed in the Customer’s cloud environment, with support across the three major cloud service providers - AWS, Azure, and GCP. This includes cross-cloud support whereby the product is deployed in one cloud and data is imported from another; and

(2) SaaS hosted – the product is hosted in Darktrace Cloud, specifically AWS. Data can be imported from Customer AWS, Azure, and GCP cloud environments.

Data that is captured, processed and analyzed is presented in the Darktrace / FAI console where users interact with the data through various dashboards and search, filtering, timeline, and file inspection capabilities.

2.2. Data retrieval and ingestion

Darktrace / FAI can retrieve and ingest data from Cloud, SaaS, and On-premise environments.

2.2.1. Sources of data

2.2.1.1. Cloud

Darktrace / FAI supports the acquisition of cloud volumes, live system data, and log-based activity. Data capture is typically achieved via APIs or export formats provided by the cloud provider. Due to the variation in APIs, authentication methods, and SDKs provided by each platform, the methodology for ingestion for each component will differ between cloud providers. The method of authentication and retrieval is outlined in the relevant documentation for each cloud provider environment, as found on the Darktrace / FAI docs website. Downloads are dependent upon the quality of network access provided by the relevant cloud provider – network access which is non-performant or unreliable may result in slow or failed downloads.

2.2.1.2. SaaS

Darktrace / FAI supports the acquisition of various SaaS log sources, including Microsoft 365, Microsoft Entra ID, and Google Workspace. The method of authentication and retrieval is outlined in the relevant documentation for each SaaS log source, as found on the Darktrace / FAI docs website. Given that data is retrieved directly from the audited activity logs of the relevant SaaS provider, returned information is limited to the events that each vendor chooses to audit and the

data recorded as part of those audit log entries. Variation in both the events available for each SaaS log source and the level of detail of the data retrieved is expected.

2.2.1.3. On-premise

Detection and alert data, as well as recorded telemetry (i.e. the real-time data collected by the detection provider, including but not limited to: processes, files, and network connections) is ingested from data retrieved from integrations with detection providers. Darktrace / FAI also supports triage acquisition of on-premise devices. The method of authentication and retrieval is outlined in the relevant documentation for each detection provider, as found on the Darktrace / FAI docs website. Typically, retrieval involves creating an API client in the detection provider with the necessary permissions for data retrieval and ingestion and subsequently adding the client ID and secret information to the Darktrace / FAI console.

2.2.2. Data capture types

Darktrace / FAI supports the following data capture types:

2.2.2.1. Full volume capture

Full volume capture involves creating a bit-by-bit copy of an entire storage device, including all allocated and unallocated sectors. This comprehensive approach enables the most complete preservation of the digital evidence, allowing for thorough analysis without altering the original data.

2.2.2.2. Triage capture

Unlike full disk acquisition, which involves making a complete copy of a storage device, triage acquisition prioritizes efficiency and speed. This acquisition method involves collecting forensic artifacts that Darktrace has determined to be the most essential, allowing for users to perform quick triage investigations. Example artifacts collected include information about running processes and network connections, contents of open files such as running binaries, file system artifacts, and system logs. Additional files and artifacts may be specified by the user to be included in the triage capture.

2.2.2.3. Log capture

Darktrace / FAI captures cloud logs from various cloud services via their API. Examples include CloudTrail and VPC Flow logs for AWS, and Activity Logs for Azure. Given that data is retrieved directly from the audited activity logs of the relevant cloud provider, returned information is limited to the events that each vendor chooses to audit and the data recorded as part of those audit log entries. Variation in both the events available for each cloud provider and the level of detail is expected.

2.2.2.4. File import

Darktrace / FAI supports importing data from cloud storage. This means users can use the Darktrace / FAI console to navigate cloud storage objects and select the relevant files to be

imported. Darktrace / FAI can process a wide range of log, file, and artifact formats as detailed below.

2.2.2.5. Supported imported data types

Currently supported data types include but are not limited to:

AWS

- EC2 (Including AMIs, EBS Snapshots, and Volumes);
- S3 Storage (both standard and Glacier);
- AWS Kubernetes (ECS and EKS);
- AWS Lambda;
- CloudTrail logs;
- GuardDuty logs;
- Kubernetes logs;
- VPC Flow logs;
- SSM logs;
- S3 Access Logs; and
- Route 53 (DNS) logs.

Azure

- Virtual Machines;
- Disks;
- Storage Containers;
- Container Blobs;
- Kubernetes (AKS);
- Activity logs; and
- Storage logs.

GCP

- Compute Engine;
- Disks;
- Storage Buckets;
- Kubernetes (GKE);
- VPC Flow logs;
- IAM logs; and
- Admin Activity logs.

Storage formats

- .E01 and split .E01;
- .VHD and .VHDX;
- .DD (Optionally .GZ Compressed);
- .GZ, .TAR, .ZIP;
- .VMDK (Partial Support); and
- Other forensic artifacts collected by EDR tools or the Darktrace / FAI triage capture tool.

OS log and artifact types – note this list is non-exhaustive

- Common AWS, Azure, and GCP logs;
- Apple System Log (ASL);
- Android usage history (app usage);
- Basic Security Module (BSM);
- Bencode files;
- Chrome Disk Cache and Preferences;
- CUPS IPP logs;

- Extensible Storage Engine (ESE) Database (EDB);
- Firefox Cache;
- Java Web Start IDX;
- JumpLists (customDestinations-ms files);
- MacOS Application Firewall, Keychain, Securityd, and Wifi logs;
- McAfee Anti-Virus logs;
- Microsoft Internet Explorer History (MSIE4-9 Cache Files or index.dat);
- Microsoft IIS logs;
- NTFS \$MFT and \$UsnJrnl:\$J;
- OLE Compound Files;
- Opera Browser History;
- OpenXML;
- Portable Executable (PE) files;
- PLSQL cache files (PL-SQL developer recall files);
- Popularity Contest logs;
- Property List (plist);
- Restore Point logs (rp.log);
- Safari Binary Cookies;
- SCCM Client logs;
- SELinux audit logs;
- SkyDrive log and error logs;
- SQLite databases;
- Symantec AV Corporate Edition and Endpoint Protection logs;
- Syslog;
- Utmp, Utmpx;
- Windows Event Logs (EVT, EVTX);
- Windows Firewall logs;
- Windows Job files (atjobs);
- Windows Prefetch files;
- Windows Recycle Bin (INFO2 and \$I/\$R);
- Windows NT Registry Files;
- Windows Shortcut Files (LNK);
- Xchat and Xchat scrollback files; and
- Zsh history files.

Volume Systems

- GPT
- LVM
- MBR
- Volume Shadow Snapshots (VSS)

File Systems

- ext2, ext3, ext4
- FAT
- NTFS (version 3)
- XFS (version 4, 5)
- Apple File System (APFS)

2.2.2.5.1. Considerations for imported data types

Split E01 disks are currently supported only in AWS. Importing zipped split E01 files is not supported. Some sub-formats of VMDK files are not supported, for better compatibility, it is recommended to convert VMDK files to .dd format before importing. Unless it is comprised of a single file, ZIP files are treated as containers of files. The supported data types listed in the previous section are subject change.

2.2.2.6. *Detections*

Darktrace / FAI supports integration with detection providers. This means as malicious activity is detected in the form of a 'detection' or 'alert' (terminology differs from vendor to vendor), Darktrace / FAI automatically performs a forensic acquisition of the impacted resources to deliver additional critical forensic level context at the host level. Given that Darktrace / FAI is ingesting detection and alert data from the detection provider, returned information is limited to the detections and alerts that each vendor chooses to surface, and thus based on the detection logic of the vendor, not that of Darktrace / FAI.

2.3. Analysis

Darktrace / FAI data analysis is multi-stage. The following outline highlights key elements of the process but is not intended to be an exhaustive list.

2.3.1.Extraction

Data is first extracted. This may involve parsing various disk formats, file systems and archive formats. At a high level, there are two data types extracted from the evidence - Events and Files. Events - Regardless of the acquisition type, as part of the evidence processing, event extraction takes place. This means that an event timeline can be constructed, consisting of file MACB events, events that have been parsed from log files, and other key events.

Files – For Customer hosted deployments only, a copy of every file that is collected is preserved. For all deployment types, the first 10KB of text of every file is also stored so that file content can be previewed in the console.

2.3.2.Detections

Static detections are run against both file and event data utilizing various detection rule formats derived from in-house and open-source rules as well as antivirus style detection capabilities. If an event or file matches the detection logic defined within the rule (for example there may be logic defined to detect one or more suspicious IP addresses, domain names, filename, process names, suspicious strings etc.), an 'alarm' is generated. The alarm can have one of three categories – (1) malicious, (2) suspicious, or (3) informational. Darktrace makes no warranty that all threats will be surfaced in the form of an alarm.

2.3.3.Automated Investigation

A machine learning model is applied to all timeline events in order to rank them. The highest ranked events represent the most significant findings identified by the automated investigation machine learning model. Events are more likely to be shown in the Automated Investigation output if they are tagged by the Darktrace / FAI alarm system or if they are similar to other alarmed events in terms of shared fields or proximity in time.

2.3.4. File analysis

Shell scripts and similar executable plain text files are analyzed using a locally run model to provide a description of functionality, if the file in question has been alarmed (see 'Detections' section above).

2.4. Output

2.4.1. Acquired data

2.4.1.1. Processed evidence

Forensic data acquired and processed from Cloud, SaaS, and On-premise environments can be investigated in a single unified view in the Darktrace / FAI console in the relevant 'Investigation'. Data can also be exported via the API and sent to webhooks for integration with third party products.

2.4.1.1.1. Detections

Detection and alert data generated from detection providers can be investigated in the Darktrace / FAI console in the 'Detections' area of the platform. Data can also be exported via the API and sent to webhooks for integration with third party products. For example, a user could create a webhook whereby any 'Detection created' event type would send detection and alert data to Jira which is used for issue tracking. Customers may also integrate findings into tools like SIEMs, task managers, and other productivity systems.

2.4.1.1.2. Response

For certain cloud resource types, users are presented with commands they may wish to execute to perform containment actions in response situations. Examples include detaching the role associated with a resource and stopping an instance.

Customer is solely liable for any impact resulting from commands or actions executed by Customer via the Darktrace / FAI console or otherwise.

2.4.2. Interface

The Darktrace / FAI console provides a series of top-level indexes that can be used to navigate and investigate the processed data.

Darktrace / FAI organizes evidence items into groups called Investigations which are displayed in the Interface. Each Investigation contains:

- An Overview of the investigation, which contains information such as an investigation summary, an overview of the attacker timeline, and key events;
- An Automated Investigation, which returns the highest ranked events representing the most significant findings identified by the machine learning model;
- Insights as to the breakdown of data and activity within the Investigation, such as determining the top 5 MITRE categories and data sources; details of the Operating System distribution of the acquired evidence; and the number of alarms generated;
- Processed evidence imported into the Investigation;

- A searchable timeline of Events, with facets to breakdown the data into separate categories; and
- Vulnerabilities that were identified within the Evidence within the Investigation.

2.4.3. Reporting

Audit logs record user activity within the platform. The audit logs can be viewed in the console or exported to CSV format. Audit actions that are recorded are as follows:

- Read;
- Create;
- Update;
- Delete;
- Login;
- Logout;
- Download;
- Reboot;
- Upgrade;
- Import;
- Upload;
- Cancel; and
- Rerun.

In practice examples include:

- A user logged into the platform;
- An investigation was created;
- An evidence item was imported; or
- An evidence item was deleted.


3. Deployment Architecture

3.1. Customer hosted

A Darktrace / FAI deployment consists of:

- **Infrastructure** – A core virtual machine that permanently runs (Forensic Acquisition & Investigation Instance' in the architecture diagram below) and temporary worker machines ('Workers') that spin up and down when processing jobs are required. A copy of every file that is collected is preserved in file storage.
- **Darktrace / FAI account/subscription/project role** – Permissions to manage the main virtual machine and start worker virtual machines if used.
- **Cross-account role** – Permissions to copy data into the main account.

Example Customer hosted architecture diagram:



Forensic Acquisition & Investigation Platform UI (Customer only) provides the ability for a user to access the Darktrace / FAI console.


Forensic Acquisition & Investigation update service (download only) provides for updating the Darktrace / FAI platform and is dependent on outbound connectivity to the Cloud Service Provider used to host the Darktrace / FAI platform. All updates to Customer hosted instances of Forensic Acquisition & Investigation are user initiated. Full details can be found on the Darktrace / FAI docs website.

Customer Threat Intelligence Sources is an optional functionality that allows Customers to integrate with various intelligence systems and incorporate custom Indicators of Compromise (IOCs) to supplement Darktrace / FAI's existing detection capabilities. For example, the Darktrace / FAI platform supports a VirusTotal API key which may be used to enhance the analysis of already detected files by comparing their file-hash with the VirusTotal database. When a file with a detection is processed, Darktrace / FAI will perform a VirusTotal lookup based on the file's hash.

Support SSH Login if required is available at the Customer's option and control. This allows the platform to be accessed via SSH in the event the UI is unavailable or unresponsive.

No Customer data is automatically sent or shared with Forensic Acquisition & Investigation online services – and sharing telemetry can be toggled on/off in the Settings area of the platform. Darktrace / FAI uses system metrics to understand the performance and feature usage patterns of the platform - e.g. type and size of evidence collected, number of projects, number of API calls, number of accounts, number of malicious and suspicious findings. No system metrics data is personally identifiable, nor does it disclose any details about investigation names or specific findings.


Darktrace / FAI can also be deployed in a cross-cloud setup, whereby the product is deployed in one Cloud Service Provider (CSP) and data is imported from another. In the architecture diagram below, this shows Darktrace / FAI deployed in AWS and data being imported from Azure and GCP.



The latest list of supported cloud regions for deploying the Darktrace / FAI product can be viewed [here](#).

3.2. SaaS hosted

In a SaaS hosted deployment, Darktrace hosts the Darktrace / FAI instance in Darktrace Cloud, specifically a multi-tenant AWS architecture. Standard Support Services apply to the SaaS hosted version of Darktrace / FAI and resultantly all Customer instances will be updated automatically.



The exact steps to access the Customer Environment via IAM authentication (in the architecture diagram above), differ between CSPs, however at a high level it involves creating a Cross-Account IAM Policy and Role in the target account/subscription/project from which the Darktrace / FAI platform will acquire data. This allows Darktrace to assume the cross-account role for data acquisition.

Darktrace / FAI comprises of cloud-hosted components and infrastructure provisioned and maintained by Darktrace Operations in the Darktrace-owned public cloud environment. All data generated is isolated within a private namespace, stored in the region chosen by the Customer in the Darktrace AWS environment. Data is encrypted at rest and Darktrace / FAI utilizes industry-standard practices to segregate tenant data.

In terms of regionality of data, the Customer will choose the AWS region they want their data to reside in. Support is currently offered for North America and Europe, with other regions to be added in the future. The default data retention period is 365 days for raw and processed evidence, as well as log and telemetry data.

4. Implementation

4.1. Customer hosted

There are two deployment types for Customer hosted – minimal and full.

A **minimal** deployment involves starting a virtual machine with only the minimal permissions required to operate. This offers a basic working environment; however, the platform will be limited in functionality. For example, the platform limits imports to triage captures and restricts simultaneous evidence processing to ensure stability. Additional instances are required to process more data or multiple evidence items concurrently. Customers that want to rapidly deploy the platform and add functionality incrementally, as well as Customers that operate highly restricted environments would benefit from the deployment type. This is available in AWS, Azure and GCP.

A **full** deployment enables the ability to utilize Workers for faster processing and importing large systems. This is available in CloudFormation in AWS, and Terraform in AWS, Azure and GCP. Once deployed, the platform creates an isolated VPC (Virtual Private Cloud) where Customers control access and permissions.

Descriptions of required permissions are available for AWS, Azure and GCP on the Darktrace / FAI docs website. Permissions are required to copy data from other accounts and cloud providers in AWS, Azure and GCP.

4.2. SaaS hosted

Darktrace will host and provision the Darktrace / FAI instance. The Customer is required to create a Darktrace / FAI role and the necessary policies for the platform to acquire evidence from the CSP account. A minimal deployment is not available in SaaS hosted deployments.

Descriptions of required permissions are available for AWS, Azure and GCP on the Darktrace / FAI docs website.

5. Administration

5.1. Settings

General administration is performed under the Settings area of the platform:

- “Detections” settings are used to configure data acquisition options for detection providers. View our documentation on how to manage Detections.
- “Intelligence” settings are used to enrich analysis. This is an optional functionality that allows Customers to integrate with various intelligence systems and incorporate custom Indicators of Compromise (IOCs) to supplement Darktrace / FAI's existing detection capabilities. For

example, the Darktrace / FAI platform supports a VirusTotal API key which may be used to enhance the analysis of already detected files by comparing their file-hash with the VirusTotal database. When a file with a detection is processed, Darktrace / FAI will perform a VirusTotal lookup based on the file's hash. More details can be found on the Darktrace / FAI docs website.

- "Processing" settings determine how the platform processes forensic evidence.
- "Preservation" settings determine how evidence and files are preserved.
- "Accounts" settings are used to connect Cloud Service Provider, and Extended Detection and Response (XDR) providers.
- "SSO" settings are used to connect oAuth and SAML providers.
- "Updates" can be used to view and trigger updates (Customer hosted deployment only).
- "Licensing" is used to view and update the license (Customer hosted deployment only).
- "API" is used to display the API reference and generate API keys.
- "SIEM" settings are used to connect a SIEM to the output of the platform.
- "Webhooks" are used to create integrations with third party platforms via webhooks.
- "Advanced" settings are used for various deployment options.
- "Experiments" are used to enable features that are currently in Beta. All Beta features are provided on an Evaluation Basis for the purposes of the Darktrace MSA and may be enabled or disabled at the Customer's discretion.
- "Profile" is used to update the current users profile information.
- "Saved queries" are used to manage saved queries available within the Event query interface.
- "Scripts" are used to manage a library of scripts that can be executed using the 'Run Command' functionality for EC2 imports.

5.2. Teams

The Teams area of the platform is where user access to both processed data and raw data stored in the cloud, as well as features, is controlled. This ensures only authorized individuals have access to specific features and data. There are five user roles in Darktrace / FAI:

- **Administrator** – Ability to access all functionality in the Darktrace / FAI platform
- **Platform Administrator** – A more restricted set of permissions than Administrator, with a focus on operational aspects of the platform including upgrades, account management, and troubleshooting.
- **Lead Analyst** – A restricted set of permissions with a focus on managing investigations including user access and taking response actions
- **Analyst** – A more restricted permission set than the Lead Analyst role, with a focus on conducting investigations including acquiring and analyzing evidence
- **Read only Analyst** – Most restricted role with read only access

A breakdown of permissions across each of the different roles and their levels of access can be found on the Darktrace / FAI docs website.

The Teams area of the platform is also used to onboard accounts from various providers so that data can be imported into the platform. For example, users would add their CSP accounts, detection provider accounts, and SaaS accounts.

6. Customer Responsibilities

The Customer is responsible for the data they import into the platform. This includes but is not limited to the sensitivity, confidentiality and legality of the data, as well as the completeness of the data. Darktrace disclaims all responsibility for incomplete or corrupt disks and or partition formats that are non-standard according to those specified in the relevant Darktrace / FAI documentation. The Customer acknowledges that the use of corrupt disks and/or non-standard partition formats may lead to parsing issues.

The Darktrace / FAI product requires specific IAM permissions to operate effectively. These are documented on the Darktrace / FAI docs website. Permissions are required to operate the main Darktrace / FAI virtual machine in the account where Darktrace / FAI is installed, and different permissions are required in target accounts to copy data back for processing. The Customer is responsible for ensuring that the requisite permissions are maintained at all times. If the Customer modifies IAM, network or other deployment configuration beyond the supported configurations, the Customer is responsible for any import failures.

The Customer is responsible for import failures and degradation of the platform when attempting to import large disks over 500GB. Such degradation may result in extended acquisition times, slower page load times, and reduced query performance.

The Customer is responsible for ensuring that an 'Account Check' is performed on all accounts that are added to the platform. This will ensure the correct permissions are assigned to be able to successfully acquire data. Where the status of any Account Check is 'Failed', the Customer is responsible for ensuring the failing permissions are addressed.

For Customer hosted deployments, the Customer is responsible for keeping the platform up to date with the most recent version or most recent prior version, and ensuring correct IAM and network access is configured. The Customer is responsible for platforming "Health-checks", which should be used periodically to confirm the platform has the correct IAM and network requirements.

7. Considerations

The Darktrace / FAI platform should be deployed with the recommended configurations – in particular:

- A supported IAM configuration that enables the required access to resources
 - o Darktrace recommends that the Darktrace / FAI platform is deployed via AWS Organizations within AWS, an Application Registration in Azure and a Service Account in Google Cloud.
 - o Darktrace / FAI should be deployed via the root account/tenant/project to sub-accounts/tenants/projects that require access.
- A supported instance specification (i.e. the core virtual machine that permanently runs, labelled as 'Forensic Acquisition & Investigation Instance' in the architecture diagram in section 2.1), defined in the CloudFormation or Terraform deployment code.
- A supported network configuration that enables access to resource endpoints, defined in the CloudFormation or Terraform deployment code.

If a Customer customizes IAM, Network and other deployment configuration beyond the supported configurations, the platform may not operate correctly.

Darktrace / FAI supports a fixed list of data types, outlined at 2.2.2.5 as well as specified in our documentation for the most up to date list. Importing unsupported data types is likely to lead to failed processing.

Darktrace / FAI analysis and monitoring is delivered by interaction with cloud provider APIs and supported integration methods. These APIs, services, and methods may incur costs within the cloud provider. Mitigations to reduce these costs are recommended where possible but are not exhaustive or unconditionally effective.

Individual vendors may place rate-limiting restrictions on the APIs and methods utilized by Darktrace / FAI. Please refer to the individual provider documentation for further information.

Darktrace / FAI retrieves data from cloud providers using the APIs and integration methods presently offered by those platforms. If this coverage or available capabilities are modified or revoked by the third-party, Darktrace will endeavor to maintain continuity but is ultimately limited in capacity to restore service.

Changes made to cloud provider capabilities or APIs are outside the scope of Darktrace control and may be introduced at any time. Changes to the known considerations due to alterations by vendors of these platforms may arise at any time.

Darktrace / FAI may be reliant upon specific vendor licenses to retrieve the events data for operation. If the required licenses are not present or are revoked, degradation or interruption of service will occur.

Where a Darktrace / FAI investigation imports significant data quantities that may significantly impact system load - "Data sampling" may be activated. Samples constitute a percentage of imported data instead of every piece of data for storage and analysis.

In the event of the Customer usage being above reasonably expected operational scope, the following limits may impact operability of the Darktrace / FAI Service:

- Maximum number of evidence items per investigation;
- Maximum number of investigations;
- Maximum number of alarms;
- Maximum number of hits per alarm;
- Maximum number of timeline events;
- Maximum size of strings stored per file;
- Enough time having elapsed before re-acquisition of a resource that has an additional alert or detection associated with it (triggered by a detection provider); or
- Auto-archiving investigations after a certain period of inactivity.

The exact details of the above listed limitations are not publicly available and are subject to change.

In addition to the limits outlined above, the following limits may also apply to the Darktrace / FAI SaaS:

- Maximum concurrent acquisitions per tenant;
- No ability to import disks >200GB; and
- Default data retention of 365 days (please note larger data retention packages can be purchased through additional SKUs).

8. Assumptions

Customer expressly acknowledges and accepts that a core functionality of the Darktrace / FAI platform is to reconstitute files recovered from Customer systems and that reconstituted files may contain harmful and/or malicious files, which if executed, could cause damage. Customer acknowledges and agrees that it is solely responsible for any and all damage caused by downloading any reconstituted files, exports, snapshots or any other data formats, whether structured or unstructured, from the Darktrace / FAI platform, however stored or accessed.

The Customer acknowledges and agrees that it is solely responsible for the lawfulness, accuracy, quality, and integrity of all Customer Data submitted to or processed through the Darktrace / FAI Service, including any special categories of personal data (as defined in Article 9 of the UK GDPR and EU GDPR) or data subject to sector-specific regulation.

Darktrace does not monitor, access, review, or validate the content of Customer Data except as strictly necessary to provide the Service in accordance with the Data Processing Addendum (DPA). Darktrace shall not be liable for any claims, losses, or damages arising from the content, accuracy, or legality of Customer Data, including where such data contains health information, financial data, or other regulated content.

The Customer warrants that it has obtained all necessary consents, authorizations, and lawful bases for the collection, processing, and transfer of Customer Data, and that such processing complies with all applicable laws, including the UK GDPR, EU GDPR, and any relevant sector-specific regulations. The Customer shall indemnify and hold harmless the Provider against any claims, fines, or liabilities arising from a breach of this clause.

Unless expressly agreed in writing, the Customer shall not upload or process through the Darktrace / FAI Service any data that is subject to specific statutory or regulatory handling requirements (including but not limited to data governed by the Payment Card Industry Data Security Standard (PCI DSS) or NHS Digital standards).

Darktrace makes no representation or warranty that the Darktrace / FAI Service will detect and generate an Alarm for all cyber threats, vulnerabilities, or malicious activity. The effectiveness of the Darktrace / FAI Service is dependent upon the accuracy, completeness, and timeliness of the Customer Data and the permissions configured by the Customer. Darktrace shall not be liable for any failure to identify threats arising from incomplete, inaccurate, low quality or incorrectly formatted Customer Data, or for any failure resulting from improperly maintained permissions and/or configurations.

Darktrace / Secure AI

- 1. Introduction 147
- 2. Product Features 148
 - 2.1. High Level Summary 148
 - 2.2. Data Retrieval and Ingestion 148
 - 2.3. Darktrace Analysis 150
 - 2.4. Output 152
 - 2.5. Interface 154
 - 2.6. Reporting 154
- 3. Deployment Architecture 155
 - 3.1. Darktrace Platform Portal 155
- 4. Implementation 155
 - 4.1. Microsoft 365 156
 - 4.2. AI Model Provider Integrations 156
 - 4.3. Shadow AI Detection 156
- 5. Administration 156
 - 5.1. User Management 156
- 6. Requirements 157
- 7. Customer Responsibilities 157
- 8. Considerations 159
- 9. Operational Change Management 160

1. Introduction

Darktrace / SECURE AI gives organizations a unified view of how AI is being used, making it possible to understand risk across prompts, agents, development environments, and unsanctioned activity.

Darktrace / SECURE AI is designed to bring clarity and control; rather than relying on predefined rules, it analyzes how prompts, agents, and AI systems behave within the context of your organization. By understanding intent across users and systems, it gives security teams the ability to see where risk is emerging and act on it in real time. The result is continuous oversight of AI usage across the business, as adoption evolves and expands.

2. Product Features

2.1. High Level Summary

Darktrace/SECURE AI extends Darktrace Self-Learning AI coverage across core prompt-based AI capabilities, low-code AI SaaS platforms, and enterprise SaaS applications with built-in AI features.

The product provides visibility and security analysis of prompt-based AI interactions, AI agent activity, and AI identity information within sanctioned enterprise AI services. Darktrace/SECURE AI also detects unsanctioned (“shadow”) AI service usage via network telemetry and third-party SASE integrations.

The system profiles users, identities, and agents based upon their activity, and of those it deems “peers” due to similar behavioral activity. Alerts are raised when specific criteria or a minimum threshold of unusual behavior is met.

2.2. Data Retrieval and Ingestion

Darktrace / SECURE AI ingests data from supported enterprise AI services, cloud and SaaS platforms, configured third-party integrations, and other Darktrace Behavioral Defense Platform components. Data types include prompt data, AI interactions, AI tool calls, audit log events (such as agent creation, permission changes, configuration modifications, and usage events), and static contextual data about users, groups, roles, permissions, and AI agents.

Retrieval and ingestion for integrated platforms is typically achieved via APIs or export formats provided by the third-party platform. Due to the variation in APIs, authentication methods, and SDKs, provided by each separate upstream vendor, the methodology for ingestion will differ between integrations and components. The method of authentication and retrieval is outlined in the relevant documentation provided for Darktrace/SECURE AI.

The scope, granularity, and timeliness of retrieved information are dependent upon the capabilities of the integrated platform, permissions granted to Darktrace, and any limitations imposed by third-party providers. Delays may be incurred where the external platform does not make events available to Darktrace for processing and analysis within the expected time frame. Delays of this nature are the responsibility of the associated third-party platform vendor.

Data ingested by Darktrace / SECURE AI is retrieved directly from the audited activity logs of the relevant third-party platforms - returned information is therefore limited to the events that each vendor chooses to audit, and the data recorded as part of those audit log entries. Variation in both the events available to Darktrace / SECURE AI, and the level of detail, are expected between platforms.

Some third-party platforms may also restrict the events and APIs available behind license requirements. Darktrace will outline where possible the minimum license requirements required for monitoring.

2.2.1. Monitored, Approved, and Shadow AI Services

Darktrace/SECURE AI distinguishes between monitored AI services, shadow AI services, and approved AI services.

- Monitored AI services are services for which Darktrace/SECURE AI has been configured to retrieve data through a supported integration or ingestion mechanism.

Shadow AI services are AI services identified through network telemetry or other detection methods which have not been approved by the customer. Operators can approve detected shadow AI services within Darktrace/SECURE AI, promoting them to approved services.

- Approved AI services are services that the customer permits for organizational use. Approval status does not imply that Darktrace/SECURE AI has visibility of activity within that service.

Darktrace/SECURE AI performs prompt analysis, activity analysis, and behavioral analysis only on data made available through supported monitoring integrations and configured data sources. Approval of a service alone does not grant visibility to Darktrace/SECURE AI.

2.2.2. Microsoft 365

Darktrace/SECURE AI integrates with enterprise software platforms such as Microsoft 365 to retrieve audited activity and AI interaction data.

Data types include prompt data from AI interactions (such as Microsoft Copilot across the Microsoft 365 ecosystem), AI agent activity and configuration data (such as Copilot Studio agents), and audit log events relating to AI usage within the platform. For real-time data such as prompts and audit events, Darktrace/SECURE AI is notified by the platform when new activity is available for retrieval.

Coverage is limited to the AI tools and interactions that the platform makes available via its supported APIs and audit mechanisms. As Microsoft 365 is continually expanding the range of tools and embedded AI experiences, not all AI-powered features or interactions within the Microsoft 365 ecosystem may be visible to Darktrace/SECURE AI at any given time. Returned information is therefore limited to what Microsoft surfaces for retrieval, and variation in both the interactions available to Darktrace/SECURE AI, and the level of detail, are expected as the platform evolves.

Darktrace/SECURE AI will also retrieve contextual enrichment data on users, agents, groups, roles, permissions, and other relevant entities from the enterprise software platform to supplement the behavioral profiles it creates.

A Microsoft 365 Global Administrator grants Darktrace/SECURE AI authorization via OAuth 2.0; using this authorization, Darktrace/SECURE AI retrieves data from the platform on an ongoing basis. Where additional platform components require separate configuration to enable monitoring (for example, Power Platform environments for Copilot Studio agent visibility), the method of authentication and retrieval is outlined in the relevant documentation provided for Darktrace/SECURE AI.

2.2.3. AI Model Provider Integrations

Darktrace/SECURE AI integrates with AI model providers to retrieve prompt activity, user-to-AI interactions, and usage patterns. Integrations are configured by the customer from the Darktrace/SECURE AI interface.

Monitoring is limited to what the configured enterprise service makes available via its APIs or other supported integration methods. The depth, format, and availability of prompt content, session information, and response data may therefore vary between providers.

Where an integration has been configured by the customer, Darktrace/SECURE AI can retrieve and ingest data regarding direct interactions with the AI model provider's service (for example, direct interactions in an enterprise ChatGPT tenant). Darktrace/SECURE AI does not automatically gain visibility into third-party applications, internal tools, SaaS platforms, custom workflows, browser-based interactions, or other systems within the customer's environment that may utilize the same AI model provider as an upstream dependency. Monitoring of such systems requires separate visibility mechanisms and is dependent upon available integrations and supported data sources.

2.2.4. Cloud Providers

Where Darktrace/CLOUD is deployed, Darktrace/SECURE AI can retrieve prompt data, audit events, and identity information from cloud providers with AI service offerings (such as AWS Bedrock). Data is retrieved via existing Darktrace/CLOUD ingestion mechanisms and routed to Darktrace/SECURE AI for analysis.

For cloud platforms, retrieved data may include AI interaction logs, agent activity, identity and permissions data, and audit events relating to AI resource creation and configuration. Darktrace/CLOUD also provides contextual enrichment data on AI agents, users, roles, permissions, and other relevant entities within the cloud environment.

This integration requires an active Darktrace/CLOUD deployment. The method of authentication and retrieval is outlined in the Darktrace/CLOUD product specification and associated documentation.

2.2.5. Shadow AI Detection

To facilitate shadow AI detection, Darktrace/SECURE AI ingests network telemetry from integrated Darktrace Behavioral Defense Platform products and/or third-party SASE integrations (such as Zscaler or Microsoft GSA) to identify connections to AI services. Network telemetry is processed to identify AI service usage patterns; raw network data is not stored beyond what is necessary to generate detections.

2.2.6. Policy Documentation

Organizational AI policy documentation can be uploaded to Darktrace/SECURE AI by an operator. Policy documents are ingested and processed to derive approved usage guidelines specific to the organization.

2.3. Darktrace Analysis

Darktrace analysis is multi-stage. The following outline highlights key elements of the process, but is not intended to be an exhaustive list of all analysis performed during operation of the Darktrace/SECURE AI product or any integrated Darktrace Behavioral Defense Platform products.

2.3.1. Prompt Analysis

Darktrace/SECURE AI processes and analyzes prompt data retrieved from third-party AI providers via a multi-stage, AI-enabled analysis pipeline.

Darktrace/SECURE AI establishes a behavioral baseline for prompt activity at the user, peer group, and organization level. This baseline is then used to identify deviations and anomalous patterns in prompt usage, and is continuously updated to ensure relevancy.

Each prompt is analyzed to derive features such as discussion topic, user intent, any potentially sensitive or malicious content, and to identify whether it constitutes an anomaly in comparison to a baseline. Prompts may also be analyzed for noncompliance with (uploaded) internal usage policies or regulatory frameworks. Where encoded or obfuscated content is identified (such as base64 encoding), Darktrace/SECURE AI will attempt to decode and analyze the underlying content. Prompts relating to software development activity are subject to the same analysis; Darktrace/SECURE AI does not perform static code analysis, code review, or vulnerability scanning of AI-generated outputs.

Darktrace/SECURE AI may also derive overall scoring for prompts for factors such as risk, harm potential, anomaly, and sensitivity. Once processed and analyzed, prompts are passed for analysis to the Darktrace/SECURE AI detection engine layer.

2.3.1.1. Session Analysis

Where session information can be reliably reconstructed from provider data, prompts are grouped into sessions for higher-level analysis. Sessions receive risk scoring, AI-generated summaries, and the aggregation of prompt-level tags and detections. Session risk scores are computed from a combination of factors including anomaly detection, presence of sensitive information, and harmful activity indicators in the constituent prompts.

Where session information cannot be reliably reconstructed, prompts may be displayed and analyzed independently.

2.3.1.2. Policy Alignment

Darktrace/SECURE AI analyzes organizational AI policy documentation to derive permitted usage rules. Prompts and other activity are then evaluated against the derived rules to highlight instances of potential policy violation.

2.3.1.3. Legislation Mapping

Darktrace/SECURE AI may also examine prompts for properties relevant to applicable legislation (such as the EU AI Act). Prompts may be tagged with categories relating to high-risk AI usage under relevant legislative frameworks.

Outputs generated by this analysis are informational indicators only. The presence of a tag, score, categorization, or alert does not indicate that a legislative violation has occurred. Likewise, the absence of a tag, score, categorization, or alert does not indicate that activity is compliant with any legislative or regulatory framework.

2.3.2. Shadow AI Analysis

Darktrace/SECURE AI analyzes network data to identify connections to AI services. Connection metadata is evaluated for indicators of AI services or AI tool interactions; where a match is identified for an unsanctioned service, a shadow AI detection is generated.

The format, method, and level of analysis performed by Darktrace/SECURE AI will be dependent upon the mode of ingestion. Analysis performed on log-level network telemetry or from pre-processed metadata supplied by SASE integrations will be subject to any limitations placed on the data issued by the third-party providers. Where full network-level visibility is configured - either supplied by existing Darktrace Behavioral Defense Platform products, or deployed to support Darktrace/SECURE AI - additional, in-depth analysis may be performed.

2.3.3. Identity Analysis

Darktrace/SECURE AI will create individualized behavioral profiles for the users and AI agents it observes. This baseline is derived from the observed activity of the individual entity and the activity of those it deems “peers” due to similar behavioral activity, and is continuously updated to ensure relevancy.

Darktrace/SECURE AI retrieves activity data for users or identities made available by third-party APIs (or other comparable methods) for analysis. Event and activity data is parsed and analyzed to identify the type of activity, any entities involved (e.g., users, files, or other resources), and to extract key information. Supplementary contextual analysis is also used to identify roles, permissions, group membership, and access scopes, which are surfaced in the Darktrace/SECURE AI interface.

2.3.4. Cross-Platform Analysis

Where other Darktrace Behavioral Defense Platform coverage areas are deployed, cross-platform analysis is performed by multiple components. For example:

- Where Darktrace/CLOUD is deployed, Darktrace/SECURE AI can retrieve and analyze audit events, identity data, and agent configuration from AI services in the cloud environment.
- Where Darktrace/EMAIL is deployed, detections relating to prompt injection attempts within inbound email may be surfaced alongside prompt-level analysis.
- Where Darktrace/NETWORK or/ENDPOINT products are deployed, Shadow AI detections are derived from network telemetry and surfaced within the Darktrace/SECURE AI interface.

The above provides illustrative examples of collaborative scenarios between platform components for reference but is non-exhaustive. Some features may also be dependent upon cross-platform components.

2.3.5. Detection Engine

The Darktrace/SECURE AI detection engine is the logical framework within which the outputs of preceding analysis are evaluated. Prompt-level analysis, session-level analysis, audit events, and their correlation are submitted to the detection engine for evaluation.

A detection scheme is defined as a set of conditions which, when met, will alert the system to the occurrence of a particular event or combination of anomalous behavior. The detection engine evaluates tags, scores, audit events, and other outputs derived during analysis to evaluate whether these criteria have been met.

Darktrace/SECURE AI provides a set of default detection schemes as standard, maintained and updated by Darktrace. Examples of conditions evaluated by default detection schemes include attempted jailbreak activity, sensitive data disclosure, possible indirect prompt injection attempts, unusual time-of-use patterns, anomalous interactions, and potential policy violations. Custom detection schemes can be defined by operators to meet specific criteria.

Care should be taken when defining custom detection schemes to ensure that the number of alerts produced is not excessive and does not impact the system's ability to function or otherwise degrade service. Similarly, modification of existing default detection schemes should not result in excessive activity or alter the logic to the extent that alerts criteria can no longer be met. The responsibility to ensure that detection schemes created or edited remain within these reasonable expected boundaries lies with the operator.

2.4. Output

Darktrace/SECURE AI surfaces the results of its analysis as alerts, scores, tags, detections, and identity insights. These outputs are accessible from the Darktrace/SECURE AI interface and, where applicable, the Darktrace Platform Portal.

2.4.1. Alerts

Darktrace/SECURE AI generates alerts for anomalous prompt or user activity, or when specific detection criteria are met. Alerts are categorized by severity (Critical, High, Medium, Trivial) and surfaced with associated risk indicators, usage context, and key detections. Examples include attempted jailbreak activity, sensitive data in prompts, possible indirect prompt injection attempts, and potential policy violations.

Alerts can be reviewed and acknowledged within the Darktrace/SECURE AI interface, or optionally synchronized to the Darktrace Platform Portal to access centralized alert management workflows. Synchronization is optional and requires customer configuration.

2.4.2. Tags

Prompts and sessions may be assigned tags. Tags summarize positive, negative, and neutral factors derived during the analysis into brief, recognizable categories for quick investigation. These include contextual information (e.g., discussion topics, derived user intent, and actions performed), and risk indicators such as policy contravention. Tags may also be applied or aggregated at the session level.

2.4.3. Scoring

Darktrace/SECURE AI produces a range of scores for prompts, sessions, identities, and user activity to capture anomaly and estimate potential risk. Multiple scores may be generated for a single entity (for example, risk, anomaly, and sensitivity), and entities may be categorized by severity (Critical, High, Medium, Trivial) based on a function of these scores.

Scores are intended to support prioritization and investigation, and do not constitute a definitive determination of malicious or non-compliant activity. The specific scores produced may also be subject to change as Darktrace/SECURE AI develops further.

2.4.4. Detected Services

Darktrace/SECURE AI will surface AI services in use within the organization. Each detected service is surfaced with associated usage trends, connected users, and key detections.

2.4.5. Shadow AI detections

Shadow AI detections identify connections to unsanctioned AI services. Each detection surfaces the AI service identified, associated user(s), and usage trends over time. Operators with appropriate permissions may mark detected services as approved.

2.4.6. Identities

Darktrace/SECURE AI surfaces users and AI agents associated with AI activity. Each identity is surfaced with contextual information such as groups, roles, permissions, and access scopes, alongside a real-time view of associated AI activity (such as prompts, sessions, and alerts).

2.4.7. Manual Response Actions

Darktrace/SECURE AI allows response actions to be manually triggered against users engaging in risky AI interactions. These actions can be created by a user with appropriate permissions from the Darktrace/SECURE AI interface.

Manual actions are automatically created in an “active” state and do not require human activation. The type of action taken (*inhibitor*) is selected during creation. The responsibility to select an appropriate inhibitor lies with the Customer.

2.4.8. Integrations

Where configured, Darktrace/SECURE AI may share alerts, risk insights, and other analysis outputs with compatible external platforms such as Microsoft Agent 365. See “2.7 Microsoft Security Ecosystem Integrations” for further details.

2.5. Interface

The Darktrace/SECURE AI is accessible from the Darktrace Platform Portal. The interface provides:

- Overview dashboards for high-priority alerts and usage trends.
- An alerts page for reviewing, investigating, and acknowledging alerts created by Darktrace/SECURE AI.
- Dedicated analysis views for prompt/session analysis and shadow AI detection.
- An identity page displaying users and AI agents with associated groups, roles, permissions, and activity.
- Configuration pages for integrations, policy management, and settings.

Darktrace/SECURE AI alerts can be optionally sent to the Platform Portal for investigation in via centralized alert management workflows.

2.6. Reporting

Darktrace/SECURE AI does not offer PDF or CSV export capabilities.

2.6.1. Microsoft Security Ecosystem Integrations

Where supported by Microsoft and configured by the Customer, Darktrace/SECURE AI may exchange high-level alert and risk information with Microsoft security and administration platforms such as Microsoft Agent 365.

Information shared through these integrations is limited to the scope defined by the integration and may include identifiers, risk indicators, alert summaries, or associated metadata. Specific data exchanged and available functionality may vary over time as Microsoft and Darktrace develop additional integration capabilities.

Use of these integrations is optional and may require explicit customer configuration or authorization.

3. Deployment Architecture

Darktrace/SECURE AI is offered in SaaS architecture model. No on-premises or client-managed deployment architecture is offered.

Darktrace/SECURE AI stores prompt data, user and agent identity information, audit logs, and associated metadata in cloud-hosted infrastructure. Prompt data, events, and alerts are retained for up to 12 weeks.

Infrastructure is provisioned and maintained by Darktrace in a Darktrace-managed public cloud environment. Darktrace/SECURE AI operates a multi-tenant architecture within regional Kubernetes clusters utilizing industry-standard best practices for data segregation. EMEA and US hosting regions are currently offered, with future regions under consideration based on demand.

Darktrace/SECURE AI infrastructure is deployed across multiple availability zones within each hosting region. Data is backed up in accordance with Darktrace operational standards to support service continuity and disaster recovery.

For in-depth Shadow AI detection, if no compatible Darktrace/HYBRID NETWORK products are present, additional deployment components such as Darktrace cSensor agents may be necessary.

3.1. Darktrace Platform Portal

Access and user management for Darktrace/SECURE AI is provided via the Darktrace Platform Portal. Darktrace/SECURE AI alerts can optionally be sent to the Platform Portal to access centralized alert management workflows.

The Platform Portal is a cloud-hosted service provided via a containerized, multi-tenant architecture. Stored data such as usernames, permissions, and actions taken in the Darktrace Platform Portal itself are globally replicated to mitigate latency.

3.1.1. Darktrace Platform Portal Gateway

Where Darktrace/HYBRID NETWORK products are integrated with Darktrace/SECURE AI to provide shadow AI detection, data is shared via the Darktrace Platform Portal Gateway service. This service allows Darktrace products to communicate through a secure, encrypted reverse tunnel proxy.

Data sent via the Platform Portal Gateway service is not stored or replicated at the Platform Portal level.

4. Implementation

Darktrace provisions the Darktrace/SECURE AI instance and ensures the customer has access to the Darktrace Platform Portal. Once provisioned, the customer connects Darktrace/SECURE AI to the enterprise platforms to be monitored. Integration configuration is performed by the customer from within the Darktrace/SECURE AI interface, accessible via the Darktrace Platform Portal.

The following outlines are illustrative and subject to change as the product and its supported integrations develop.

4.1. Microsoft 365

A Microsoft 365 Global Administrator authenticates Darktrace/SECURE AI via OAuth 2.0 from within the Darktrace/SECURE AI interface, and consents to the permissions requested by Darktrace. Audit logging must be enabled in the Microsoft tenant intended for monitoring.

For Power Platform monitoring (Copilot Studio agent visibility), an application user must be created for Darktrace/SECURE AI in each Power Platform environment intended for monitoring. This step requires a user with the System Administrator role in the relevant environment.

4.2. AI Model Provider Integrations

For AI model provider integrations, the customer configures the integration from the Darktrace/SECURE AI Integrations page. Configuration typically requires an API key or compliance access key with appropriate permissions, generated by an administrator of the relevant platform. Specific requirements vary by provider and are outlined in the relevant product documentation.

4.3. Shadow AI Detection

Where shadow AI detection requires integration with existing Darktrace/HYBRID NETWORK products, a connection via the Darktrace Platform Portal Gateway is required. Where no compatible Darktrace/HYBRID NETWORK products are present, additional deployment components (such as Darktrace cSensor agents) may be necessary; implementation of these components is performed by Darktrace in line with the relevant product's implementation process.

Detailed deployment steps for each integration are outlined in the relevant product documentation.

5. Administration

Within the Darktrace/SECURE AI interface:

- Operators may add, remove, or re-authenticate integrations with enterprise AI services.
- Operators may upload organizational AI usage policies and configure which policy rules are active.
- Operators with the relevant permissions may approve or reject AI services identified by Shadow AI Detection.

Access to the underlying systems of Darktrace/SECURE AI is reserved for authorized Darktrace personnel only.

5.1. User Management

User access and permissions for Darktrace/SECURE AI are managed via the Darktrace Platform Portal. Darktrace/SECURE AI does not provide its own user administration capabilities.

Role-based permissions control what each user can view and act upon within Darktrace/SECURE AI. Permissions that grant access to raw prompt data should be assigned with due care given the sensitive nature of the underlying content.

6. Requirements

The following requirements are illustrative and subject to change as the product and its supported integrations develop.

- Access to the Darktrace Platform Portal (for Darktrace/SECURE AI interface access, and user management)/
- For Microsoft 365 monitoring:
 - A Microsoft 365 Global Administrator account to complete the OAuth 2.0 authorization process.
 - Audit logging enabled in the Microsoft tenant intended for monitoring.
 - For Power Platform monitoring (optional), additional configuration must be performed by an account with the System Administrator role in each Power Platform environment intended for monitoring.
- For AI model provider integrations:
 - An enterprise-tier plan or equivalent plan with the relevant AI model provider that provides the necessary access to retrieve prompt or activity data (specific tier requirements vary by provider).
 - An API key or compliance access key with appropriate permissions, generated by an administrator or owner of the relevant platform.
- For shadow AI detection:
 - For SASE-based detection: an active integration with a supported SASE provider (such as Zscaler or Microsoft GSA).
 - For network-based detection via existing Darktrace/HYBRID NETWORK products: the Darktrace instance must be connected to the Platform Portal, and internet access must be enabled for the Platform Portal module.
 - Where no compatible Darktrace/HYBRID NETWORK products are present, additional deployment components (such as Darktrace cSensor agents) may be required.
- For cloud AI platform monitoring:
 - An active Darktrace/CLOUD deployment for the relevant cloud environment.

7. Customer Responsibilities

Notwithstanding any additional responsibilities set out in any other section of this Product Specification, and without limitation, Customer's responsibilities are to ensure that:

- Darktrace/SECURE AI integrations have been configured and authorized according to Darktrace instructions.

This includes ensuring that all permissions requested Darktrace/SECURE AI are granted, whether during an authentication flow such as OAuth, granted separately as permissions or roles in the third-party platform, or granted/ provided in another comparable deployment process.

- Authentication for integrations is performed by - or granted via - a user with the appropriate permissions as outlined in the relevant Darktrace/SECURE AI documentation.

Recommendations or requirements placed on authorizing users are due to technical necessities. Darktrace cannot guarantee that OAuth grants, API keys, or other authentication methods performed by users without the appropriate permissions will result in functional operation.

- Darktrace/SECURE AI retains authentication to integrated platforms and, if lost, is re-authenticated in a timely manner.

If a key is revoked, expires, or is otherwise invalidated, Darktrace/SECURE AI will cease to receive data from that integration until a valid key is re-configured.

- Darktrace/SECURE AI integrations are re-authenticated when requested by Darktrace. Darktrace/SECURE AI may require re-authentication to add additional capabilities or to support upstream changes in vendor APIs.
- Modifications are not made to the third-party platform configuration or licensing that would prevent operation.

Examples include the removal of auditing within the platform, removal of licenses required to contact specific APIs, removal of permissions from associated API keys, or any other configuration alteration that would interrupt service operation.

- End-user access is managed appropriately and within recommended best practice.
 - Any permissions or roles that grant visibility of raw prompt data should be assigned with due care, and with due diligence and data security considerations in mind.
- Changes to the approval status of AI services in shadow AI detection accurately reflect the organization's current posture. If a service is marked as approved and the organization's posture subsequently changes, the Customer is responsible for updating this status.
- Organizational AI policies uploaded to Darktrace/SECURE AI must be valid, current, and kept actively updated.
- Where organizational AI policies reference supporting policies, governance frameworks, standards, procedures, or other documentation, all relevant documents are uploaded and maintained within Darktrace/SECURE AI. Darktrace/SECURE AI can only derive rules from policy information made available to it. References to external or unavailable documents may result in incomplete policy interpretation or policy-alignment analysis.
- Integrations and data sources from which Darktrace/SECURE AI retrieves information are configured and maintained appropriately. Darktrace is not responsible for limitations in the scope, completeness, or granularity of data made available by third-party platforms, and analysis is limited to the data received through configured integrations.
- Changes made to components or overall system configuration by users of the Darktrace platform do not impact the system's ability to function or otherwise degrade service.
 - Examples of user error which may result in degraded service include the exclusion of key users using activity filters, the exclusion of all users from Darktrace Autonomous Response, the creation of excessively overactive detection schemes, and the modification of existing detection schemes to the extent that alerts criteria can no longer be met.

8. Considerations

Delays may be incurred where the external platform does not make prompts, events, or other interaction data available to Darktrace/SECURE AI for processing and analysis within the expected time frame. Delays of this nature are the responsibility of the associated third-party platform vendor. Individual vendors may also place rate-limiting restrictions on the APIs and methods utilized by Darktrace/SECURE AI. Maximum prompt length constraints may also apply.

Darktrace/SECURE AI monitors and retrieves data from third-party platforms using the APIs and integration methods presently offered by those platforms. If this coverage or available capabilities are modified or revoked by the third-party, Darktrace will endeavor to maintain continuity but is ultimately limited in capacity to restore service. Changes made to third-party platform capabilities or APIs are outside the scope of Darktrace control and may be introduced at any time. Changes to the known considerations due to alterations by vendors of these platforms may also arise at any time.

Response actions are subject to the restrictions placed by third-party platforms on API capabilities. This can result in specific user types being ineligible for actions, specific circumstances impacting the efficacy of actions, and the requirement for additional configuration (in both Darktrace and/or the third-party platform). Individual considerations are outlined in product documentation where known.

Darktrace/SECURE AI may be reliant upon specific vendor licenses to retrieve the events necessary for operation or to trigger responses. If the required licenses are not present, or are revoked, degradation or interruption of service will occur.

The scope, granularity, and timeliness of retrieved information are dependent upon the capabilities of the integrated platform, permissions granted to Darktrace, and any limitations imposed by third-party providers. Data ingested by Darktrace / SECURE AI is retrieved directly from the audited activity logs of the relevant third-party platforms - returned information is therefore limited to the events that each vendor chooses to audit, and the data recorded as part of those audit log entries.

Some AI services may not expose complete prompt visibility, session visibility, identity context, workflow information, or audit information through supported APIs. Visibility of agent-to-agent communication, orchestration activity, internal workflow execution, tool delegation, and multi-agent interactions may also be limited or unavailable depending on the provider. Variation in both the events available to Darktrace / SECURE AI, and the level of detail, is expected between platforms.

Session reconstruction may not be possible where providers do not expose sufficient information for accurate session correlation. In such cases, Darktrace may use available information to provide best-effort session analysis. Third-party providers may also release new AI features or capabilities for which audit logging or prompt data is not immediately available via their APIs. Darktrace/SECURE AI is unable to monitor interactions that are not surfaced by the third-party provider.

For integrated services, Darktrace/SECURE AI can only retrieve and ingest data regarding direct interactions with the AI model provider's service. Darktrace/SECURE AI cannot retrieve interactions with tools or SaaS services within the customer's environment that may themselves use the AI model provider as an upstream vendor. Similarly, Darktrace/SECURE AI cannot retrieve prompt data for interactions with AI model providers performed by employees through unsanctioned channels (for example, interactions conducted in private browsing sessions or via non-enterprise user accounts).

Darktrace / SECURE AI supports multilingual prompt analysis. Analysis quality may vary by language, translation quality, language-detection accuracy, and provider capabilities. Analysis of non-text content, uploaded files, images, and other non-textual AI inputs may also be limited depending upon provider visibility and supported functionality.

Evaluation of prompts against organizational policies or legislative frameworks is intended to provide informational indicators only. Outputs from this analysis do not constitute a guarantee of compliance with any legislation or organizational policy. Darktrace/SECURE AI may not detect all instances of non-

compliance, and detections are not guaranteed to be accurate in all circumstances. Darktrace/SECURE AI will not enforce organizational policy rules which are discriminatory, illegal, or may be deemed in contravention of relevant legislation. Darktrace/SECURE AI may require acceptance of usage terms prior to policy upload.

Shadow AI visibility is not expected to be exhaustive. Detection may be limited by the completeness, quality, or depth of available network data, or by unobservable or deliberately obfuscated usage patterns.

Darktrace/SECURE AI cannot be integrated with Darktrace Behavioral Defense Platform components to support shadow AI detection without a connection via the Darktrace Platform Portal. Additionally, if advanced shadow AI detection requires additional Darktrace/HYBRID NETWORK deployment components, all considerations, requirements, and customer responsibilities also apply to that component of the deployment.

Darktrace/SECURE AI utilizes AI service offerings from cloud providers in the course of its analysis. Data residency for processing is determined by the mechanisms and service regions of the upstream cloud provider.

Where alert synchronization to the Darktrace Platform Portal is enabled, customers should consider any applicable data residency, regulatory, or organizational requirements.

Darktrace/SECURE AI generates operational and service-health metrics to support capacity management, service reliability, and product improvement. Such metrics may include aggregated utilization information, performance measurements, detection statistics, and other metadata necessary for operation of the service. These operational metrics are necessary for service delivery and cannot be individually disabled.

Darktrace/SECURE AI utilizes generative AI technology. Generative outputs may be subject to error.

9. Operational Change Management

Darktrace/SECURE AI is delivered as a SaaS service. Updates, improvements, and maintenance are applied by Darktrace without requiring customer action.

Where Darktrace/SECURE AI integrates with other Darktrace Behavioral Defense Platform products (such as Darktrace/HYBRID NETWORK or Darktrace/ENDPOINT), version upgrades to those products may introduce changes that affect cross-platform functionality, available data sources, or integration behavior.

Further detail is provided in the product-agnostic section on “Operational Change Management and Version Updates”.

Product Agnostic Details

1. Cyber AI Analyst.....	161
1.1. Cyber AI Analyst Hypotheses	162
2. Documentation	163
3. Operational Change Management and Version Updates	163
3.1. Operational Change Management	163
3.2. Version Updates	163
3.2.1. Model Updates	164
4. Data Backups and Retention.....	164
5. Appliances	166
5.1. Safety Information for Darktrace physical instances	167
5.2. Ownership and Return.....	167
5.3. Delivery	167
6. Assumptions.....	168

1. Cyber AI Analyst

Darktrace models are used as a trigger to invoke Cyber AI Analyst. When the conditions for a model are met, a model alert is created; Cyber AI Analyst reviews and investigates all relevant model alerts that occur on the system as a starting point for its analysis process. The output from this analysis process is Cyber AI Analyst Incidents - a collection of one or more related events of anomalous activity. Incidents are formed through a meta-analysis of activity type, entity type (such as devices, identities), and endpoints involved in each event. Each incident can encompass multiple stages of activity as it develops.

The Darktrace Cyber AI Analyst operates a hypothesis-based analysis approach, where activity is evaluated against a number of possible, relevant hypotheses and a determination is taken of which (if any) hold based upon the evidence gathered and investigations performed. This investigation process involves numerous forms of data analysis including, but not limited to, AI and Machine Learning algorithmic approaches, statistical analysis techniques, and other forms of natural language and mathematical analysis.

The Cyber AI Analyst will combine activity across different Darktrace Real-Time Detection coverage areas where possible.

Although a model alert may be the trigger for an investigation, that does not mean the activity Cyber AI Analyst surfaces is directly related to the original model alert. The behavioral analysis it performs may discover anomalies or patterns of activity that were not the original trigger point for the model alert but are worthy of investigation. Similarly, very few model alerts that are investigated will result in an incident - only activity the Cyber AI Analyst considers high priority. Whether Cyber AI Analyst has created a related incident is displayed in appropriate locations within the Darktrace Threat Visualizer and Darktrace Mobile App.

Users may manually trigger Cyber AI Analyst investigations into devices of interest or trigger Cyber AI Analyst investigations through third-party telemetry inputs.

1.1. Cyber AI Analyst Hypotheses

As of Darktrace 6.2, users may alter the hypotheses investigated by Darktrace Cyber AI Analyst as a result of a model alert or add hypotheses for investigation to custom models.

Darktrace strongly recommends that operators do not alter the hypotheses associated with Darktrace-maintained models; any activity which is not investigated due to a modification of this type is the responsibility of the Customer.

Darktrace also strongly recommends limiting the number of custom models Cyber AI Analyst is invoked by; a significant number of additional investigations will create additional load on the Cyber AI Analyst engine, potentially resulting in failure to investigate key activity.

1.2. Darktrace aianalyst.darktrace.com

aianalyst.darktrace.com is a new bi-directional data delivery channel for event data, diagnostic data, and global threat intelligence.

It operates alongside Darktrace Call Home. The “AI Analyst Connectivity” module on the Darktrace System Config page contains settings for this new service, including whether the Darktrace ActiveAI Security Platform is opted-in to receive threat intelligence packages (“Darktrace Inoculation”) or opted-in globally.

The service is enabled by default for Call Home customers on upgrade to Darktrace 6.2. However, this service is not enabled on deployments with a non-standard or region-restricted Call Home configuration.

The aianalyst.darktrace.com service is used for rapid data delivery between the Darktrace ActiveAI Security Platform and a centralized Darktrace location. The data sent via this channel includes contextual external hostname data (rarity), and a subset of that currently transmitted by Darktrace Call Home such as alerts, event metadata, and system health data.

The service has four primary purposes:

- Allow the Cyber AI Analyst located locally on the Darktrace ActiveAI Security platform to access global threat intelligence rapidly and on-demand (inbound).
- Send health and diagnostics data to Darktrace operations more frequently than over Darktrace Call Home (outbound).
- Send event metadata back to Darktrace for per-deployment monitoring (outbound).
- Deliver the Darktrace Inoculation threat intelligence feed (inbound).

When enabled, Darktrace Cyber AI Analyst will request additional data from the service to supplement and enhance the investigations it performs. This information includes global hostname rarity and other metadata. As the service operates on a “push” basis, Darktrace support services can respond to alerts and critical events faster than those received via Call Home.

1.2.1. Inoculation

Darktrace Inoculation is a threat intelligence feed available to all clients with the aianalyst.darktrace.com service.

When an emerging threat is identified by Cyber AI Analyst, anonymized metadata from its investigation findings are sent back to a central location. This metadata is used to construct a Darktrace Inoculation package - a collection of behavioral indicators which can be used by the Darktrace real-time detection engine to identify similar threats in new environments. The package is distributed via the aianalyst.darktrace.com service to all Darktrace ActiveAI Security Platform instances with “Inoculation” enabled.

1.2.2. aianalyst.darktrace.com and Call Home

The aianalyst.darktrace.com service is provided alongside Darktrace Call Home as an additional data delivery channel. The service is enabled by default for Call Home customers.

Data sent via aianalyst.darktrace.com is a subset of that already monitored via the Darktrace Call Home service.

Data sent via the aianalyst.darktrace.com is done so on a “push” basis, ensuring near real-time delivery.

Darktrace Call Home instead operates a regular polling methodology for data delivery and retrieval. Unlike Darktrace Call Home, the aianalyst.darktrace.com service does not provide any mechanism for remote access

2. Documentation

Documentation for the installation, operation and administration of Darktrace environments and Products is provided on the Darktrace Customer Portal.

Customer has the right to make a commercially reasonable number of copies of the Documentation, provided however, that Customer must reproduce and include all of Darktrace's and its suppliers' copyright notices and proprietary legends on each such copy.

3. Operational Change Management and Version Updates

3.1. Operational Change Management

Darktrace follows a set of standardized methodologies and procedures for efficient and prompt handling of changes to infrastructure and application, which enables beneficial changes to be made with minimal disruption to the service. In the event of reasonably anticipated or actual sustained unavailability of hosted infrastructure beyond Darktrace's reasonable control, to enable Darktrace to enact its disaster recovery and business continuity practices and maintain continuity of the provision of the Offering, Darktrace may take appropriate and proportionate measures to relocate Customer Data to another hosted location.

All Darktrace software is developed under secure development policies and practices. This is applicable to software run on Darktrace supplied hardware, managed by Darktrace operations in private cloud environments, or operated on third-party hardware (for example, Darktrace / ENDPOINT cSensor agents).

3.2. Version Updates

Version updates to Darktrace / NETWORK and Darktrace / IDENTITY are supplied through Threat Visualizer software bundles.

When “Call-Home” is enabled, all Darktrace Master appliances will automatically be upgraded by Darktrace to the latest Threat Visualizer release, unless an approval requirement has been configured and registered with Darktrace operations. Where possible, updates will be applied outside standard business hours. If this is not possible, the update process will cause minimal disruption for Threat Visualizer users.

If approval is required, an upgrade can be manually initiated from the management console or by Darktrace operations.

Where “Call-Home” is not enabled, software bundles are available from the Darktrace Customer Portal and can be unpacked and installed within the management console. Configuration of both automatic and manual updates in the management console is described in the relevant product documentation.

Darktrace vSensors will automatically update if granted access to the Darktrace update packages infrastructure. The update schedule for the Darktrace osSensor and other Agents offered in container format is defined by the client during configuration.

For cloud-hosted environments, software updates are managed by Darktrace operations. Threat Visualizer software is automatically updated when a new version becomes available; where possible, updates will be applied outside standard business hours. If this is not possible, the update process will cause minimal disruption for Threat Visualizer users.

Software releases are subject to the Darktrace Product Support and End Of Life (EOL) Policy, which can be found in the Customer Portal.

3.2.1. Model Updates

Darktrace will periodically update the standard supplied Darktrace Real-Time Detection models - customers with Call-Home or cloud-hosted instances will receive updates automatically, clients without automatic updates will receive all applicable model updates when Darktrace Threat Visualizer software is updated. The conditions for auto-update eligibility are described in the relevant product documentation.

4. Data Backups and Retention

For cloud-hosted instances managed by Darktrace operations, multiple short-term snapshot backups are taken on a rolling basis to ensure continuity in a disaster recovery scenario.

For physical appliances, Darktrace provides mechanisms to create backups of physical appliance configuration which can be exported to an external server. Backups should be treated as sensitive data and access should be protected. Backups can be created either manually, or automatically each day. By default, only the three most recent backups will be retained in the Darktrace appliance's local storage. If a new backup file is created, the oldest backup in local storage will be discarded. A backup will include all Darktrace machine learning, models, and model alerts, as well as subnet information, device information, and configuration settings on the Threat Visualizer. A backup will *not* include transactional data such as connections in the Event Log, Advanced Search entries, and PCAP files, nor will it include configuration settings on the console menu.

For Darktrace deployments with Probe and Master appliances, only the Master appliance needs to be backed up. For Unified View deployments, or deployments with more than one Master appliance, all Master appliances must be backed up individually.

It may be necessary to re-authenticate a Darktrace / IDENTITY module after a restoration from backup, depending upon the specific module and deployment configuration.

Data retention for metadata output and combined log data is on a rolling basis and is dependent upon a number of factors such as hardware capability, traffic makeup and other operational components which contribute data to the platform. Transactional data has a typical data retention of around 30 days.

5. Appliances

The technical specifications of each physical appliance size are detailed below:

	DCIP-S	DCIP-M	DCIP-X2	DCIP-Z	DCIP-XA
Form factor	1U rack mountable (Half-depth)	1U rack mountable	2U Rack mountable	2U rack mountable	2U rack mountable
Dimensions (in)	17.32" x 14.57" x 1.73"	17.32" x 29.33" x 1.73"	17.32" x 29.33" x 1.73"	17.32" x 29.33" x 1.73"	17.32" x 29.33" x 1.73"
Dimensions (cm)	44 x 37 x 4.4	45 x 74.5 x 4.4	46 x 74.5 x 4.4	47 x 74.5 x 4.4	48 x 74.5 x 4.4
Weight (lbs / Kg)	13.3lbs / 6kg	33lbs / 15kg	51lbs / 23kg	51lbs / 23kg	51lbs / 23kg
Racking	19" rack	19" rack	19" rack	19" rack	19" rack
Admin Interface	1x10/100/1000 BASE-T	1x1000 BASE-T	1x1000 BASE-T	1x1000 BASE-T	1x1000 BASE-T
Remote Management Interface	1x10/100/1000 BASE-T	1x1000 BASE-T	1x1000 BASE-T	1x1000 BASE-T	1x1000 BASE-T
Copper analysis interphases	3x10/100/1000 BASE-T	3x1000 BASE-T	1x1000 BASE-T, 2x 10G BASE-T	1x1000 BASE-T, 2x 10G BASE-T	N/A
Fibre analysis interfaces	N/A	2x10Gbe/1Gbe SFP+	2x10Gbe/1Gbe SFP+	2x10Gbe/1Gbe SFP+	4x10Gbe/1Gbe SFP+ OR 1x40Gbe QSFP+ on FPGA NIC
Power Supply	Single 350W IEC 13C 100/240V	Dual 1100W IEC 13C 100/240V	Dual 1300W IEC 13C 100/240V	Dual 1300W IEC 13C 100/240V	Dual 1300W IEC 13C 100/240V
Power Consumption	Idle: 26W - 89 BTU/hr	Idle: 120W - 409 BTU/hr	Idle: 128W - 436 BTU/hr	Idle: 128W - 436 BTU/hr	Idle: 128W - 436 BTU/hr
	85%: 89W - 305 BTU/hr	85%: 359W - 1224 BTU/hr	85%: 365W - 1245 BTU/hr	85%: 365W - 1245 BTU/hr	85%: 365W - 1245 BTU/hr
	Max: 105W - 358 BTU/hr	Max: 418W - 1426 BTU/hr	Max: 426W - 1453 BTU/hr	Max: 426W - 1453 BTU/hr	Max: 426W - 1453 BTU/hr
Supported Expansion Modules	Can support one expansion model: 2-port 1G/10G SFP+ 2-port 1G RJ45 1000 BASE-T 4-port 1G RJ45 1000 BASE-T	Can support one expansion model: 2-port 1G/10G SFP+ 2-port 10G RJ45 10000 BASE-T 2-port 1G RJ45 1000 BASE-T 4-port 1G RJ45 1000 BASE-T	Can support up to three expansion models: 2-port 1G/10G SFP+ 2-port 10G RJ45 10000 BASE-T 2-port 1G RJ45 1000 BASE-T 4-port 1G RJ45 1000 BASE-T	Can support up to three expansion models: 2-port 1G/10G SFP+ 2-port 10G RJ45 10000 BASE-T 2-port 1G RJ45 1000 BASE-T 4-port 1G RJ45 1000 BASE-T	N/A
Safety Certificate	UL 60950-CSA 60950, EN 60950, IEC 60950 CB Certificate & Report, IEC 60950				
EMI Certification	FCC Part 15, Class A (CFR 47) (USA), ICES-003 Class A				

5.1. Safety Information for Darktrace physical instances

Darktrace appliances must operate within thermal limits, such that the ambient inlet temperature must never exceed 35°C (95°F). Continued operation close to this limit may impair the long-term reliability of the system. Appliances are intended to operate in environments that meet ASHRAE Class A2 data center guidelines. A specialized rugged hardware Probe is available for industrial environments unsuitable for standard hardware.

All drive bays must be occupied. Empty drive bays should be occupied with a blank drive supplied by Darktrace.

Darktrace appliances are provided with a physical security seal on the chassis and a front bezel to protect the drive bays. The cover should not be removed by anyone other than a Darktrace engineer or at a minimum, under the remote supervision of a Darktrace engineer. When a security seal must be removed, new seals will be provided and should be applied immediately after work requiring the seal to be removed is complete.

5.2. Ownership and Return

Unless otherwise agreed to in writing, title to all Appliances (and all components thereof) provided by Darktrace to Customer will always remain with Darktrace. Customer's use of any Appliance is subject to this Product Specification.

Upon termination or expiration of the Evaluation Period or Subscription Period (as applicable), Customer shall:

- (a) promptly return all Appliances to Darktrace or the applicable Partner, in accordance with Darktrace's (or the applicable Partner's) instructions; and
- (b) ensure all Customer Data is removed from the Appliance. Darktrace shall not be responsible for maintaining or protecting any configuration setting or data found on the returned Appliance or component part of the Appliance and it is Customer's sole responsibility to delete any such information prior to return.

5.3. Delivery

Darktrace uses commercially reasonable efforts to ship Appliances on the delivery dates agreed in writing by Darktrace and the Customer (or Partner, if applicable); provided, however, that Customer's sole and exclusive remedy for any delay in delivery or for failure to give notice of delay shall be for Darktrace to make such delivery as soon as practicable thereafter.

Darktrace is not liable for the acts or omissions of any third-party courier or shipping provider. Darktrace may withhold or delay shipment of any order if Customer, or the applicable Partner, is late in payment or is otherwise in default under the Agreement or the Partner Arrangement.

The Appliance is provided solely as the medium for delivery and operation of the Software and must not be used for any other purpose. Whilst the Appliance is in Customer's possession, Customer must: (i) store and use the Appliance in a proper manner in conditions which adequately protect and preserve the Appliance; (ii) not sell, charge, pledge, mortgage or otherwise dispose of the Appliance or any part of it; (iii) not permit any lien to arise over the Appliance (or part thereof); and (iv) keep the Appliance free from distress, execution and other legal process.

Customer shall be responsible for preparing the delivery location for the delivery of the Appliance(s) and for the provision of all necessary access and facilities reasonably required to deliver and install the Appliance(s). If Darktrace is prevented from carrying out delivery or installation of any Appliance because

no such preparation has been carried out, Darktrace may levy additional charges to recover its loss arising from such event.

In the event that there are additional fees attributed to delivering to the Customer's designated delivery site, Darktrace shall deliver the Appliance(s) FCA (Incoterms 2010) to the agreed sites and unless otherwise set out in the applicable Product Order Form, Customer shall pay and be exclusively liable for all costs associated with shipping and delivery including without limitation, freight, shipping, customs charges and expenses, cost of special packaging or handling and insurance premiums incurred by Darktrace in connection with the shipment of the Appliance(s). Unless otherwise set out in the applicable Product Order Form or agreed in writing by Darktrace and the Customer (or Partner, if applicable), Darktrace will ship by the method of its choice.

Darktrace shall identify itself in all documents related to the shipment of the Appliance(s) as the exporter of record from the applicable jurisdiction of export, and Customer (or its agent, as applicable) as the importer of record into the country of delivery. Darktrace can provide Customer with reasonable assistance and support relating to the import of the Appliances.

6. Assumptions

The Software may contain or be accompanied by certain Third- Party Products including Open-Source Software. Any Open-Source Software provided to Customer as part of the Offering is copyrighted and is made available to Customer under the GPL/LGPL and other Open-Source Software licenses. Copies of, or references to, those licenses may be set out in a Product Order Form, the Third-Party Product packaging or in a text file, installation file or folder accompanying the Software. If delivery of Open Source Software source code is required by the applicable license, Customer may obtain the complete corresponding Open Source Software source code for a period of three years after Darktrace's last shipment of the Software by sending a request to: Attn: Legal Department - Open Source Software Request, Darktrace Holdings Limited, Maurice Wilkes Building, Cowley Road, Cambridge CB4 0DS, United Kingdom. All other implied licenses are disclaimed, and all rights not expressly granted herein are reserved to Darktrace.

Unless expressly agreed between the Parties in writing, the Offering does not include the monitoring, interpretation or corrective action with respect to any Alerts. No advice, report, or information, whether oral or written, obtained by Customer from Darktrace or through or from the use of the Offering shall create any warranty not expressly stated in this Agreement. Customer understands that: (a) any outcome of the use of the Offering involving security assessment is limited to a point-in-time examination of Customer's security status; and (b) the Offering does not constitute any form of representation, warranty or guarantee that Customer's systems are secure from every form of attack, even if fully implemented. Customer understands and acknowledges that not all anomalies / intrusions may be reported or prevented.

The Customer shall perform or procure the performance of the responsibilities set out in this Service Definition in a manner which ensures Darktrace is not delayed from performing its obligations in accordance with this Agreement. To the extent that the Customer's delay or failure to comply with a responsibility does or may cause Darktrace to miss any timeframe for the performance of an obligation, Darktrace shall be entitled to an extension of time equivalent to the delay caused by the failure of the Customer.

Customer will own all rights, titles and interests in and to the Customer Data and the contents of any Alerts. In respect of any Customer Data stored on the Appliance, Customer grants to Darktrace a limited and non-exclusive license to access and use the Customer Data only to the extent necessary for Darktrace to perform the Services. Customer agrees Darktrace may utilize the details of any Alerts evaluations occurring in Customer's network and any connected data source on an anonymized basis and excluding any Customer Confidential Information and / or Personal Data, to develop and improve the Darktrace technology. Customer is solely responsible for its use of the Offering, the activities of its users

and for the accuracy, integrity, legality, reliability and appropriateness of all Customer Data. Customer expressly recognizes that Darktrace does not create nor endorse any Customer Data processed by or used in conjunction with the Offering. Customer further acknowledges that Darktrace and its Affiliates do not provide or undertake backup or maintenance services for Customer Data and Customer undertakes that it shall be solely responsible for backup of all Customer Data.

Upon expiry of the Subscription, Darktrace shall maintain Customer Data and grant Customer access to the cloud Services, solely to download and delete any Customer Data. Thereafter, Darktrace will delete or destroy all copies of Customer Data without liability or additional notice, unless legally prohibited from doing so. Customer Data cannot be recovered once deleted or destroyed.